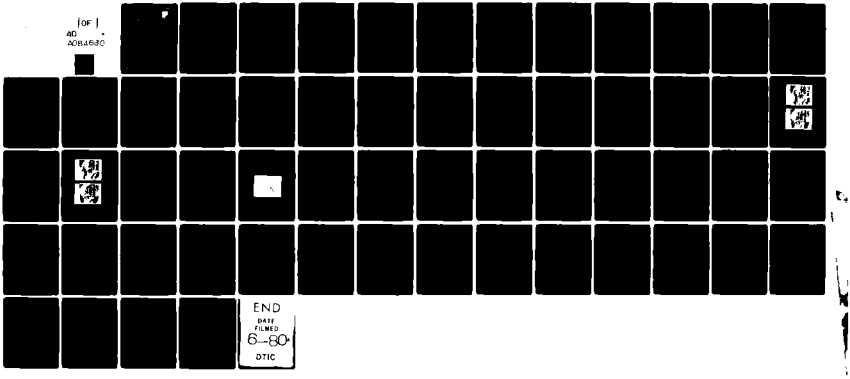
AD-A084 630 LOCKHEED MISSILES AND SPACE CO INC PALO ALTO CA PALO --ETC F/G 17/9
SAR OBJECT CHANGE DETECTION STUDY.(U)
MAR 80 K DUTTA, C D KUGLIN F30602-78-C-0347
UNCLASSIFIED LMSC/D676644 RADC -TR-80-69 NL



RADC-TR-80-69
Final Technical Report
March 1980

LEVEL II

12

AD A084630

SAR OBJECT CHANGE DETECTION STUDY

Lockheed Missiles and Space Company

Kalyan Dutta
Charles D. Kuglin

DTIC
ELECTE
MAY 27 1980

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

ROME AIR DEVELOPMENT CENTER
Air Force Systems Command
Griffiss Air Force Base, New York 13441

DDC FILE COPY

80 5 27 005

This report has been reviewed by the Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RADC-TR-80-69 has been reviewed and is approved for publication.

APPROVED: Douglas J Praska

DOUGLAS J. PRASKA, 2Lt, USAF
Project Engineer

APPROVED: Howard Davis

HOWARD DAVIS
Technical Director
Intelligence & Reconnaissance Division

FOR THE COMMANDER: John P. Huss

JOHN P. HUSS
Acting Chief, Plans Office

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (IRRE), Griffiss AFB NY 13441. This will assist us in maintaining a current mailing list.

Do not return this copy. Retain or destroy.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: RADC-TR-80-69	2. GOVT ACCESSION NO.: AD-A084630
3. RECIPIENT'S CATALOG NUMBER:	
4. TITLE (and Subtitle): SAR OBJECT CHANGE DETECTION STUDY.	5. TYPE OF REPORT & PERIOD COVERED: Final Technical Report. 29 Sep 78 — 15 Jun 79.
6. PERFORMING ORG. REPORT NUMBER: LMSC-D676644	
7. AUTHOR(s): Kalyan Dutta, Charles D. Kuglin	8. CONTRACT OR GRANT NUMBER(s): F30602-78-C-0347
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Lockheed Missiles & Space Company, Lockheed Palo Alto Research Laboratory 3251 Hanover St Palo Alto CA 94304	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 62702F 45940050
11. CONTROLLING OFFICE NAME AND ADDRESS: Rome Air Development Center (IRRE) Griffiss AFB NY 13441	12. REPORT DATE: March 1980
	13. NUMBER OF PAGES: 58
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): Same	15. SECURITY CLASS. (of this report): UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE: N/A

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

Same

18. SUPPLEMENTARY NOTES

RADC Project Engineer: Douglas J. Praska, 2Lt (IRRE)

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Synthetic Aperture Radar	Phase Correlation
Change Detection	Fourier Spectrum
Scene Descriptors	Moment Invariants
Image Registration	

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A study has been performed to assess the applicability of three area-based change-detection algorithms to the problem of recognition of differences between two images. These approaches are based on defining sets of descriptors for small regions of each scene, and then comparing these sets of descriptors to indicate areas of change. The first algorithm is based on the phase-correlation concept, which is a technique for extracting accurate estimates of the relative displacement of two

(Cont'd)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

210118

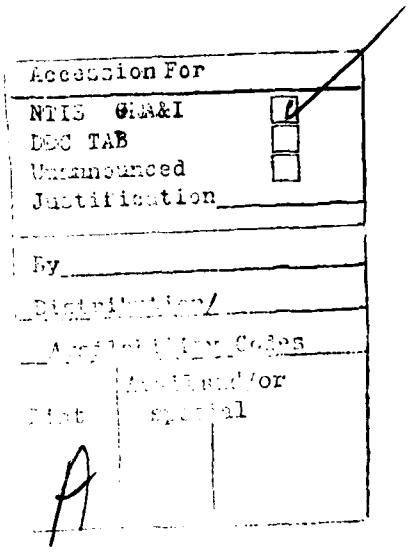
Item 20 (Cont'd)

images from the phase of their cross-power spectrum. The ~~second~~ is based on using a subset of power-spectral coefficients for each image. ~~This~~ subset is chosen ~~so as~~ to minimize differences due to scene bias, shading, or high-frequency noise content. The ~~third technique~~ is based on defining image moments and a set of functions of these moments which are invariant with respect to a number of geometrical distortion or scale changes of the image.

These algorithms ~~have been~~ applied to furnished SAR images. The results of the three algorithms are consistent among themselves, but correlate weakly with known ground truth information. Noise analysis of the furnished images indicates that they are excessively noisy. Studies to determine the size of targets that can be reliably detected by use of these methods indicate that the size of targets in the furnished images fails to meet the size criterion for reliable change detection. It is concluded that, for imagery of this quality and resolution, ~~that~~ target detection and association algorithms should prove to be more effective for change detection.

PREFACE

This report, SAR Object Change Detection, was prepared by Lockheed Missiles & Space Company, Inc. (LMSC) for the Rome Air Development Center (RADC), Griffiss Air Force Base under Contract F30602-78-C-0347. The study was conducted from 29 September 1978 to 15 June 1979. The objectives have been to assess the applicability of three region-based change-detection methods to synthetic aperture radar imagery.



CONTENTS

Section		Page
1.	INTRODUCTION AND SUMMARY	1-1
	1.1 Background	1-1
	1.2 Use of Region-Based Descriptors for Automatic Change Detection	1-1
	1.3 Summary of Results	1-2
2.	TECHNICAL APPROACH	2-1
	2.1 Change Detection Based on Phase Correlation	2-1
	2.2 Change Detection Based on Power Spectral Coefficients	2-4
	2.3 Change Detection Based on Image Moment Invariants	2-6
3.	ANALYSIS OF SAR IMAGES USING REGION-BASED DESCRIPTORS	3-1
	3.1 SAR Image Description	3-1
	3.2 Image Preprocessing	3-2
	3.3 Image Quality Evaluation Based on Available Ground-Truth Information	3-4
	3.4 Evaluation of Image Quality by Statistical Analysis	3-6
	3.5 Image Segmentation Scheme	3-13
	3.6 Data Processing Sequence	3-16
	3.7 Data Analysis	3-19
	3.8 Combination of Algorithms	3-20
	3.9 Detection Rate and False-Alarm Estimates	3-24
4.	SENSITIVITY ANALYSIS	4-1
	4.1 Image Overlap	4-1
	4.2 Target Size	4-1
	4.3 Image Signal-to-Noise Ratio	4-5
5.	REFERENCES	5-1

ILLUSTRATIONS

Figure		Page
2-1	The Phase-Correlation Algorithm	2-2
2-2	Displacement Error and Probability of False-Match Contours	2-3
2-3	Phase Correlation Algorithm Applied to Change Detection	2-5
2-4	Power-Spectrum Correlation Algorithm Applied To Change Detection	2-7
2-5	The Use of Moment Invariants Descriptors in Change Detection	2-8
3-1	Representative SAR Imagery for Change Detection	3-3
3-2	Estimated Locations of Targets in Situation 1 and Situation 2 Images	3-5
3-3	Full-Resolution 512 x 512 Pixel Segment of SAR Imagery for Target Visibility Experiment	3-8
3-4	Image Segmentation Scheme and Processing Sequence	3-15
3-5	Effective Signal/Noise Ratio as a Function of Lack of Overlap Between Two Image Segments for Invariant Descriptors	3-17
3-6	Theoretical Matrix of Target Change; SAR Images From Situations 1 and 2	3-21
3-7	Phase Correlation Change Detection Matrix; SAR Images From Situations 1 and 2	3-22
3-8	Regions of Change Labeled by Phase Correlation Algorithm	3-25
3-9	Regions of Change Labeled by Power Spectrum Correlation Method	3-26
3-10	Regions of Change Labeled by Invariant Moments Method	3-27
4-1	Phase-Correlation Method – Cosine Weighting Used; Calculated Correlation Between Two 32 x 32 Image Segments of SAPPHIRE SAR Image, Full-Resolution, 49 Segments Averaged	4-2
4-2	Power-Spectrum Correlation Method – Cosine Weighting Used; Calculated Correlation Between Two 32 x 32 Image Segments of SAPPHIRE SAR Image, Full-Resolution. 49 Segments Averaged	4-3
4-3	Invariants-Correlation Method – Cosine Weighting Used; Calculated Correlation Between Two 32 x 32 Image Segments of SAPPHIRE SAR Image, Full-Resolution. 49 Segments Averaged	

Figure

Page

4-4 Results of Embedding an Artificial Target of Intensity 256. Background Maximum: 230 4-6

TABLES

Table		Page
2-1	Functions of Central Moments Which Are Invariant Under Translation and Rotation	2-9
2-2	Functions Invariant Under Translation, Rotation, Contrast Change, and Scale Change	2-10
3-1	SAPPHIRE SAR Image Data Used for Change Detection	3-1
3-2	Summary of Visual-Static Target Detections in Pixel-Averaged SAR-Image Frame	3-7
3-3	Signal-To-Noise Ratio Estimates for Full-Resolution SAR Images	3-11
3-4	Signal-To-Noise Ratio Estimates for Reduced-Resolution SAR Images	3-12
3-5	Signal/Noise Ratio Estimates for Two Identical SAR Image Segments to Which A Controlled Amount of Noise Has Been Added	3-13
3-6	Results of Correlation Between the Various Change-Detection Techniques	3-23

EVALUATION

The work performed in this effort contributed towards the goals of Technical Program Objective (TPO) R2C. Current automatic change detection techniques require high subpixel alignment accuracies and have unacceptably high false alarm rates. This effort examined three unique, region-based, automatic change detection techniques as a means of investigating new potential methods which do not require high subpixel alignment accuracies and which exhibit acceptable false alarm rates. The results of this study demonstrate the characteristcs, limitations and capabilities of region based techniques when applied to Synthetic Aperature Radar (SAR) imagery.

Douglas J. Praska

DOUGLAS J. PRASKA, 2LT, USAF
Project Engineer

Section 1

INTRODUCTION AND SUMMARY

1.1 BACKGROUND

The problem of recognition of differences between two images is of interest in a wide range of applications. An important example is the analysis of side-looking radar imagery, for which images taken at various times must be compared with a reference image in order to discern areas where significant differences in detail exist. The problem is made difficult because of the large volumes of data which must be inspected in order to detect changes. If human photointerpretation facilities are not to be overloaded, the process of change detection must be made automatic, at least in the initial stages, the attention of the human interpreter being directed only to those few areas where significant changes are likely to have occurred.

Presently developed methods of change detection rely on a combination of image alignment between the reference and mission images, removal of relative distortions, grey-shade compensation, and image subtraction in order to detect areas of change. In order to improve the accuracy of change detection, LMSC proposed an approach based on defining sets of descriptors for small regions of each scene, and then comparing these sets of descriptors to indicate areas of change. These ideas were directly relevant to some of the SAR change-detection problems of concern to the RADC, and a short study was performed to develop our approach and to conduct preliminary evaluations of these methods, using data furnished by RADC. This report summarizes the results obtained from that study.

1.2 USE OF REGION-BASED DESCRIPTORS FOR AUTOMATIC CHANGE DETECTION

Present change detection methods which use image alignment and subtraction suffer from a number of disadvantages. Accuracy of alignment is usually critical, and unless subpixel accuracy can be obtained, the number of false alarms is usually unacceptably high. Such precision of alignment is extremely difficult to achieve, especially when the imagery has been acquired at different perspectives and when there is significant terrain relief.

These methods would be desirable if changes on the order of one pixel were of some significance to a photointerpreter. Usually, however, objects of interest are represented by a number of pixels, so that pixel-by-pixel comparisons may not be required, or even desirable, because of the attendant false-alarm problem. Thus, rather than determining changes on a pixel-by-pixel basis, what is required is that regions of change be detected, with the detection process being insensitive to local differences in contrast or noise.

This view led to the LMSC developed approaches, in which the basic idea is to define a set of descriptors on small regions of each scene to be compared. The values of these descriptors are then compared to indicate regions of change. These methods all avoid the need for extremely precise preregistration of image frames.

In addition, since the number of descriptors is usually much smaller than the number of pixels in each region, a certain amount of data compression is also achieved using region-based descriptors.

In the change-detection approach proposed by LMSC, three kinds of region-based descriptors were described:

Phase Correlation. The phase-correlation method (Refs. 1, 2, and 3) determines the amount by which two similar scenes are displaced from each other; this is done by computing the inverse Fourier transform of the phase terms of their cross-power spectrum. In applying this technique to the problem of change detection, the phase-correlation profile is computed over small regions of contiguous pixels in each pair of images. A strong peak near the origin indicates that the images being compared are sufficiently alike. Conversely, a peak which is displaced from the origin indicates that features in one image are displaced with respect to features in the other image, and the absence of a strong peak indicates a lack of similarity between the two scenes. The last two cases can be used to indicate a region of change between the two images.

Power-Spectral Correlation. For two image segments which are sufficiently alike, the power-spectral coefficients, as computed by Fourier transforming the images, are also similar. Moreover, for features which are merely shifted from one image to the other, the power spectra are identical, since all information regarding the shift resides only in the phase of the Fourier spectrum. By retaining only a subset of the power-spectral coefficients which are likely to contain information about changes of interest, some data reduction is possible using this method. In applying power-spectral correlation to change detection, the power spectra are again computed over small regions of contiguous pixels in each image after the images have been approximately aligned. A subset of the spectral coefficients are then correlated, and a sufficiently low correlation value indicates a region of change in the image.

Moment-Invariant Functions. In the third method proposed, descriptors are defined in the spatial domain, and consist of various functions defined on moments of the pair of images. As before, small regions of each image are considered, and image moments are computed. A small set of moment functions is then calculated, and the degree of correlation between the moment functions of two image regions indicates the degree of similarity between these regions.

Each of these proposed methods is detailed in Section 2. All of these methods were implemented in this study, and applied to the data furnished. The details of the analysis are provided in Section 3; the results are discussed in Section 5. A summary of the results is given in the following subsection.

1.3 SUMMARY OF RESULTS

SAR data furnished by RADC was preprocessed and aligned approximately, prior to change detection analysis. Each of the three proposed change-detection algorithms was applied to the SAR images, and matrices describing the degree of change between pairs of SAR images were obtained. The results of performing a correlation among the three sets of matrices indicate that, while the results from the three change-detection methods are consistent, they are only weakly correlated

with a change matrix calculated on the basis of known, ground-truth data. A particular combination of the change-detection methods which was used also correlated very weakly with the known matrix of target changes.

Because of these conclusions, a set of experiments aimed at determining the quality of the available SAR images was devised. The SAR images were subjected to a noise-analysis procedure and statistical measures of the estimated signal-to-noise ratios in the images were developed. The signal-to-noise ratios were measured. The results of the statistical analyses indicated that the available SAR images are excessively noisy.

The images were also analyzed from a knowledge of the available ground-truth information. The analysis showed that an unacceptable number of targets of interest are not visible in the furnished SAR images, thus giving rise to considerable differences between ground truth and image truth. The primary reason for this difference appears to be the presence of large areas of low SAR signal return in the images. Because of the presence of these shadow areas, the images themselves are poorly correlated with the furnished, ground-truth information.

Based on the results obtained from the analysis of noise and from the analysis of target visibility, we concluded that the usefulness of the three area-based change-detection methods we investigated cannot be judged from the results obtained using this particular set of data. For these change-detection methods to be effective, the signal-to-noise ratio in the available images must substantially exceed unity. It was also clear that for reliable change detection a minimum target size is necessary, and that most of the targets in the furnished, SAR images fail to meet the size criterion for reliable change detection.

We conclude that, if data of the quality available for this preliminary evaluation is to be processed by an automatic change-detection system, then certain other approaches, based on a target-association scheme, should perhaps be used instead. In one such scheme currently under investigation by LMSC, a method called symbolic matching with confidence evaluation is used. This method extracts certain high-level information from each image, such as target shape, size, orientation and spatial relationships to other targets. Targets and changes in target configuration are detected by analyzing this information. For each detected target, a confidence level is assigned, and for each confidence level, an associated list of targets can be generated. Reference and mission images are analyzed in this way to generate target lists, and these lists are compared to obtain the change-detection result. The advantages of this approach are: a high-degree of data reduction, potential for real-time processing, greater tolerances for imprecise alignment of scenes, and variations in target return.

Section 2

TECHNICAL APPROACH

2.1 CHANGE DETECTION BASED ON PHASE CORRELATION

The LMSC-developed phase-correlation algorithm is a technique for extracting an accurate estimate of the relative displacement between two images from the phase of their cross-power spectrum (Refs. 1, 2, and 3). The phase-correlation function is obtained by first computing the cross-power spectrum as is done in the case of cross correlation. The phase of the cross spectrum is then extracted at each spatial frequency by the operation

$$e^{j\phi} = \frac{G_1 \, G_2^*}{|G_1 \, G_2^*|} \qquad (2\text{-}1)$$

and the inverse transform

$$d = F^{-1}\left[e^{j\phi} \right] \qquad (2\text{-}2)$$

is then computed. Here G_1 and G_2 are the two-dimensional, forward Fourier transforms, and F^{-1} denotes the inverse Fourier-transform operator. Figure 2-1 shows a block diagram of the phase-correlation extraction procedure.

This procedure can be shown to yield all information about the translation between two images; at the same time, the "whitening" of the cross spectrum, resulting from ignoring the amplitude spectrum, can be shown to remove most scene dependence and to yield a narrow peak whose location is directly related to the vector translation between the images. Also, as in the case of normalized cross correlation, the phase-correlation peak location and amplitude is invariant under a scaling or a level shift of the image-intensity function.

The peak height of the primary peak obtained from the phase-correlation algorithm procedure can be shown to be a sensitive indicator of scene similarity (Ref. 2). Given the number of pixels used, the probability of a false match is determined only by the peak height in a manner shown in Figure 2.2. Scene changes can be detected by performing a series of correlation computations over local regions of two images which are only approximately aligned, and by comparing the peak amplitudes obtained. Regions of change produce peak heights which are significantly lower than average.

In the present study, the similarity measure used in connection with the phase correlation algorithm is one which considers both the primary-peak amplitude and its displacement away from the origin. For the purposes of change detection, it is desirable to indicate a region of change either if the computed-peak amplitude is sufficiently low, or if the primary peak is sufficiently displaced from the origin. Therefore, a similarity measure which takes into account peak amplitude, as well as

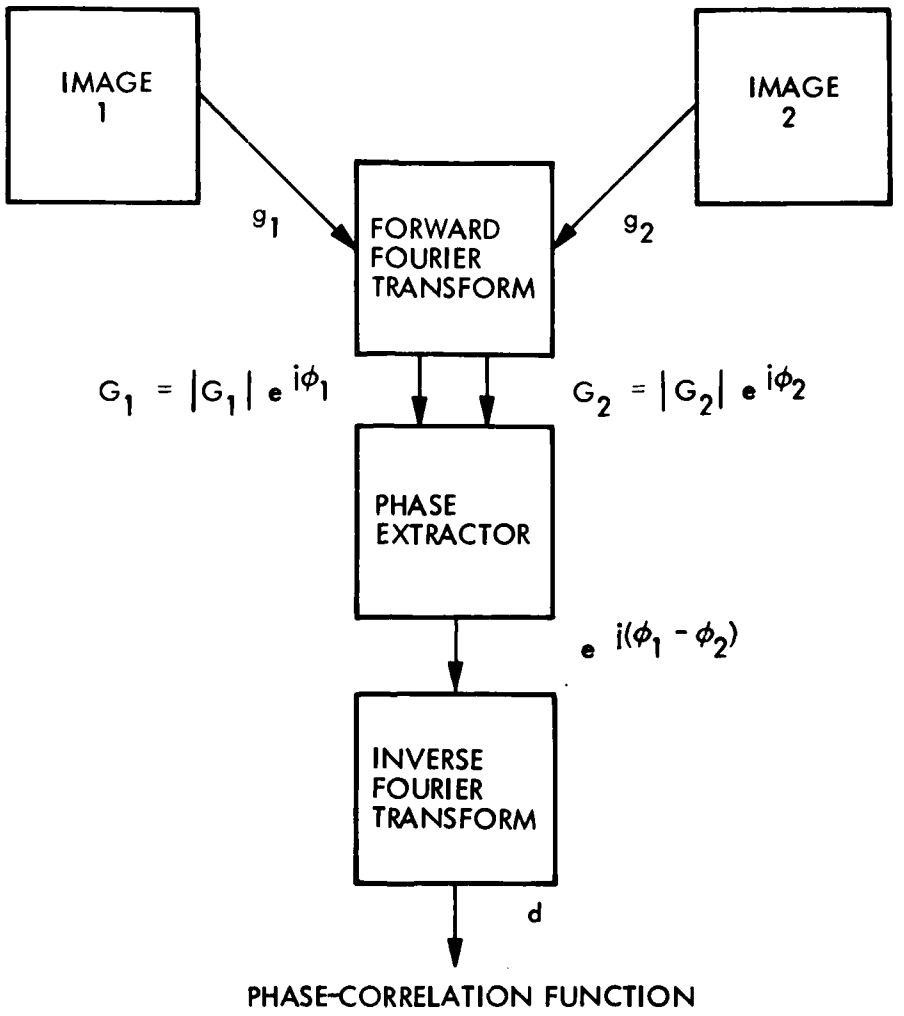


Figure 2-1 The Phase-Correlation Algorithm

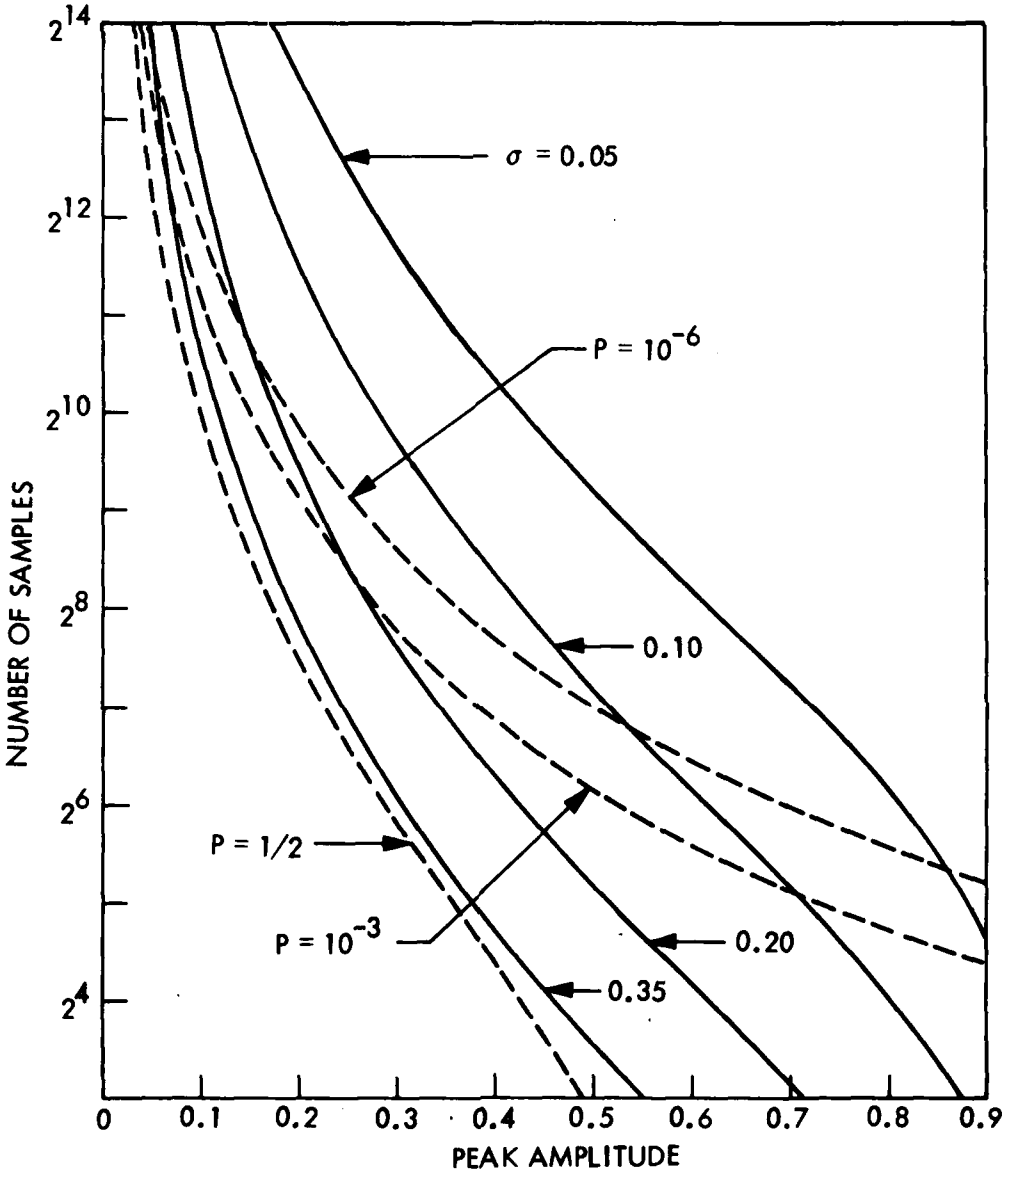


Figure 2-2 Displacement Error and Probability of False-Match Contours

diplacement, was constructed. If P denotes the primary-peak amplitude, and $\vec{L}$ denotes the computed vector displacement of this peak, then a definition of the similarity measure is

$$S = \frac{P}{1 + \sigma|\vec{L}|^2} \quad . \tag{2-3}$$

Here, σ is a conveniently chosen scale factor which determines the relative importance given to peak displacement from the origin.

Computation of the similarity measure is shown in diagrammatic form in Figure 2-3. In applying the phase correlation algorithm to the problem of change detection,it is sufficient to determine the peak amplitude, and a displacement vector for each pair of regions in the two images. This information can be combined later in a variety of ways to construct the most appropriate form of similarity measure for a particular situation.

Unlike the other two methods, change detection by the phase-correlation technique does not allow the possibility of data reduction, because the correlation between each region of the reference and mission images must be computed. Thus, the entire set of pixels for each image must be stored and used in the computation.

Finally, regions of probable change are flagged by analyzing the ensemble of similarity measures generated for the set of elementary regions. Regions for which the similarity measure fails to exceed a certain, fixed threshold can be tagged as regions of change. Or, a variable threshold can be computed, so a fixed fraction of elementary regions is always selected for possible change.

2.2 CHANGE DETECTION BASED ON POWER-SPECTRAL COEFFICIENTS

The Fourier power spectrum of an image, being the transform of the image auto-correlation, is insensitive to translations of the image. In the case of the discrete Fourier transform, this statement is strictly true only for images composed of objects located within a uniform intensity border, or for cyclically shifted images. The spectra of finite, continuous images will approximately satisfy this property, provided the translations are small. Thus the power spectra of two similar images will be quite similar, even though the images are only approximately aligned. Therefore spectrum-based descriptors can be used as a basis for change detection.

In the present study, emphasis is placed on those regions of the spectrum which are most likely to contain information about changes of interest. Since spatial frequencies in the neighborhood of the origin contain primarily information about image bias and shading,a certain region around the origin can be ignored. Also, since the highest-spatial frequencies are likely to be noise,primarily if the image has not been adequately sampled, spatial frequencies above a certain cutoff value may also be eliminated. This procedure allows a certain amount of data reduction in defining a set of image descriptors for each elementary region of the images to be compared. In addition, the fact that the Fourier power spectrum has the property of inversion symmetry is useful in reducing the data store requirement by an additional factor of 2.

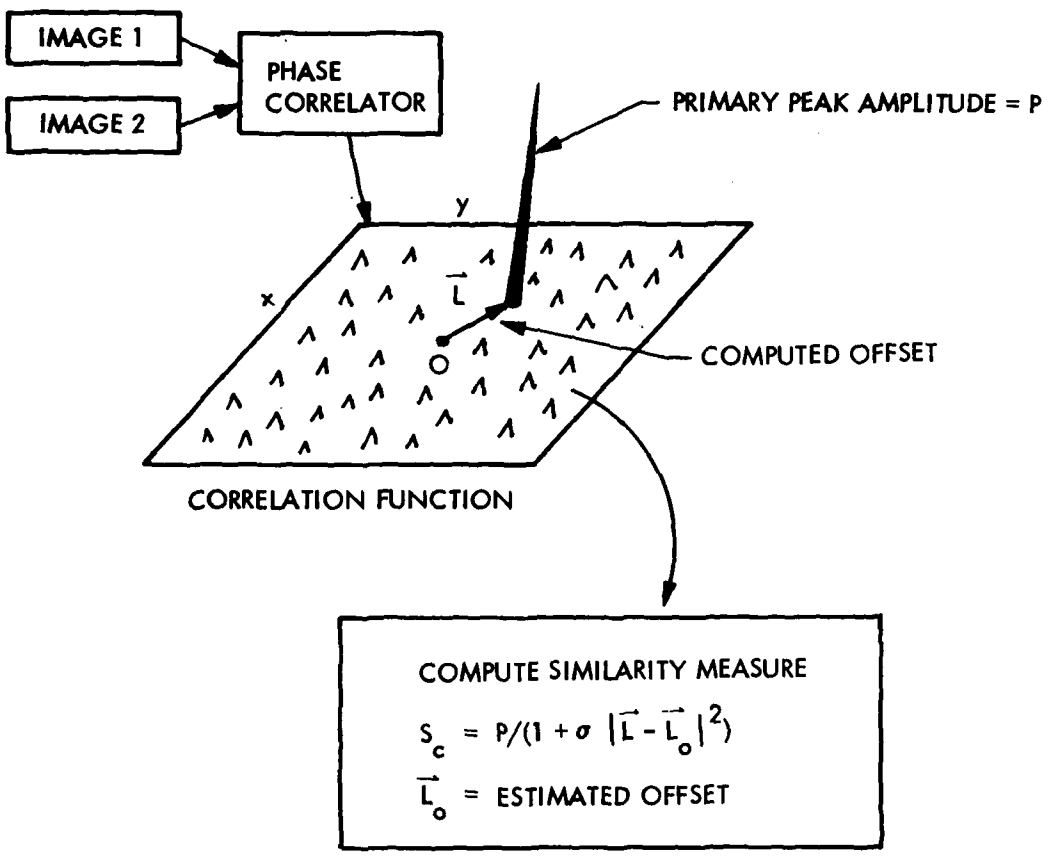


Figure 2-3 Phase-Correlation Algorithm Applied to Change Detection

Figure 2-4 shows the power-spectral coefficient procedure as applied to the change detection problem. An elementary region in the image is assumed; for this region, the Fourier spectrum is calculated, then a certain annular segment of these coefficients is retained, as the set of image descriptors for that elementary region. One-half of the spectral coefficients in the annulus are redundant and need not be stored.

In determining the degree of similarity between two images, first the images are approximately aligned, and an identical set of elementary regions is defined on each. The above subset of power-spectral coefficients is then calculated for each elementary region. Two corresponding elementary regions are determined to be identical if the correlation coefficient between the two sets of spectral descriptors is sufficiently large, that is, the function f of Figure 2-4 is simply the correlation coefficient. The use of such a procedure eliminates the effects of simple contrast differences or other scaling effects between the two images.

As in the case of phase correlation, the ensemble of similarity measures (in this case the correlation coefficient) can now be analyzed and regions of change flagged. In this second method, the important advantage is that the reference-image pixels are not required to be stored. Only a small fraction of the power-spectral coefficients for each region are retained.

2.3 CHANGE DETECTION BASED ON IMAGE MOMENT INVARIANTS

The third method which has been investigated for change detection is based on the set of descriptors referred to as invariants of image moments. Given an elementary region of an image, it is possible to define a set of quantities that can be shown to be invariant with respect to a number of transformations of the original image (Ref. 4 and 5). Since the moment invariants method has not been so widely used as, for instance, the phase-correlation method, some of its details are given here.

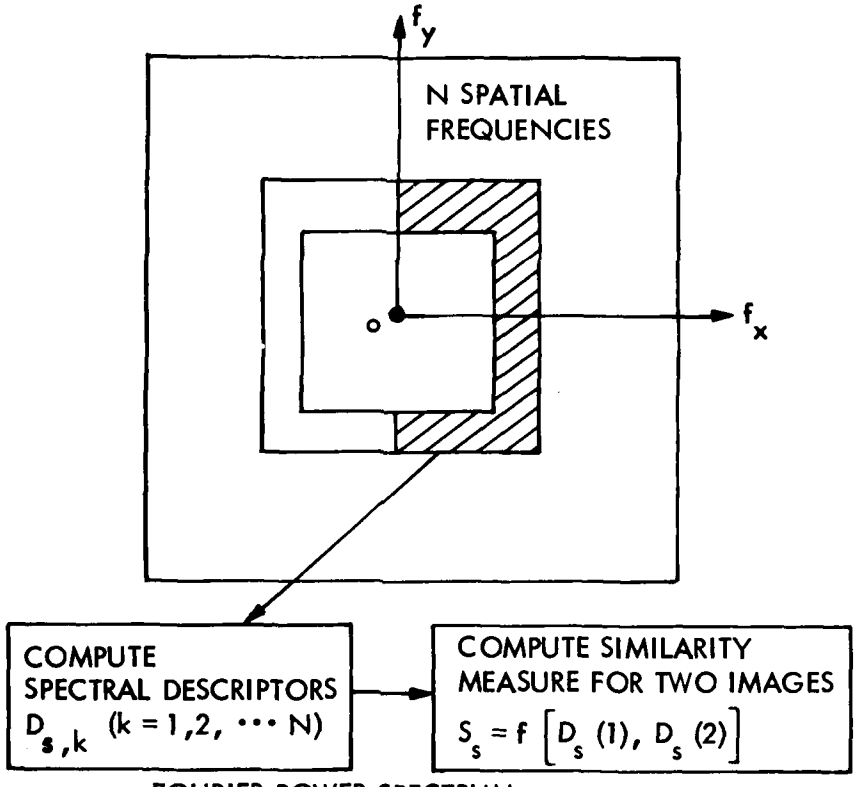
For a two-dimensional image $f(x, y)$, a set of moments can be defined, as

$$m_{pq} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} x^p \, y^q \, f(x, y) \, dx \, dy \quad , \quad p, q = 1, 2, \text{---}. \tag{2-4}$$

A modified set of moments which are insensitive to a translation of the image can be defined as:

$$\mu_{pq} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} (x-\bar{x})^p \, (y-\bar{y})^q \, f(x, y) \, dx \, dy \tag{2-5}$$

$$\bar{x} = m_{10}/m_{00} \quad , \quad \bar{y} = m_{01}/m_{00} .$$



FOURIER POWER SPECTRUM:

$$P(\vec{f}) = |G(\vec{f})|^2$$

$$G(\vec{f}) = F\{g(\vec{r})\}$$

INVERSION SYMMETRY:

$$P(\vec{f}) = P(-\vec{f})$$

Figure 2-4 Power-Spectrum Correlation Algorithm Applied to Change Detection

These are referred to as the central moments of the image (referring to the image centroid).

Based on the set of central moments μ_{pq}, it is possible to define (Ref. 4) a set of seven invariant functions $\phi(n)$, $n = 1, --7$. These functions (the definitions of which are given in Table 2-1), can be shown to be insensitive to translation as well as to rotation of the original image. A further set of six quantities $\beta(N)$, $N = 1, --6$, can be defined such that they are invariant to rotation, translation, scaling changes, and contrast changes in the original image. These quantities, defined in terms of the ϕ functions, are listed in Table 2-2 (see Ref. 5).

As in the case of the previous two methods, the above described invariances are, of course, only true for images embedded in a uniform background; for finite, continuous scenes, some variation will occur. Again, extremely precise alignment of image frames to be compared is not needed, and only a very small number of descriptors (6) are required to be stored for each elementary region in the reference image. Therefore, this technique has the greatest potential for data reduction. As before, change detection is accomplished by comparing the sets of moment invariants β defined oncorresponding, elementary regions. In the present study, a correlation-coefficient function was used for the comparison.

Figure 2-5 shows the way in which the moment-invariants method is applied to the problem of change detection.

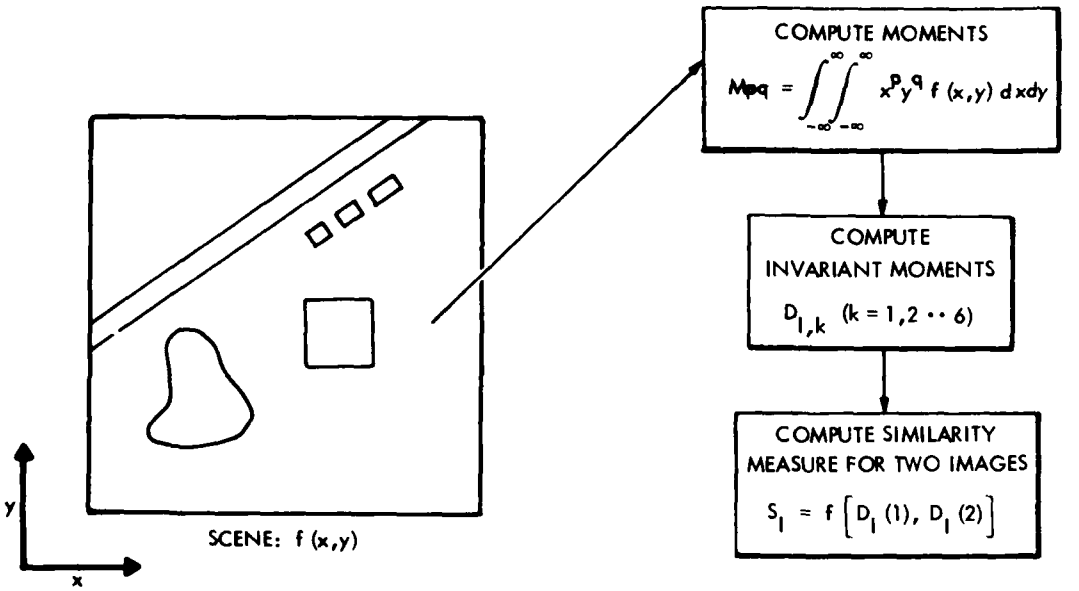


Figure 2-5 The Use of Moment-Invariants Descriptors in Change Detection

Table 2-1

FUNCTIONS OF CENTRAL MOMENTS WHICH ARE INVARIANT UNDER TRANSLATION AND ROTATION

$$\phi(1) = \mu_{20} + \mu_{02}$$

$$\phi(2) = (\mu_{20} - \mu_{02})^2 + 4\mu_{11}^2$$

$$\phi(3) = (\mu_{30} - 3\mu_{12})^2 + (3\mu_{21} - \mu_{03})^2$$

$$\phi(4) = (\mu_{30} - \mu_{12})^2 + (\mu_{21} + \mu_{03})^2$$

$$\phi(5) = (\mu_{30} - 3\mu_{12})(\mu_{30} + \mu_{12})\left[(\mu_{30} + \mu_{12})^2 - 3(\mu_{21} + \mu_{03})^2\right] + (3\mu_{21} - \mu_{03})(\mu_{21} + \mu_{03})\left[3(\mu_{30} + \mu_{12})^2 - (\mu_{21} + \mu_{03})^2\right]$$

$$\phi(6) = (\mu_{20} - \mu_{02})\left[(\mu_{30} + \mu_{12})^2 - (\mu_{21} + \mu_{03})^2\right] + 4\mu_{11}(\mu_{30} + \mu_{12})(\mu_{21} + \mu_{03})$$

$$\phi(7) = (3\mu_{21} - \mu_{03})(\mu_{30} + \mu_{12})\left[(\mu_{30} + \mu_{12})^2 - 3(\mu_{21} + \mu_{03})^2\right] - (\mu_{30} - 3\mu_{12})(\mu_{21} + \mu_{03})\left[3(\mu_{30} + \mu_{12})^2 - (\mu_{21} + \mu_{03})^2\right]$$

Table 2-2

FUNCTIONS INVARIANT UNDER TRANSLATION, ROTATION, CONTRAST CHANGE, AND SCALE CHANGE

$$\beta(1) = \frac{\sqrt{\phi(2)}}{\phi(1)} \quad , \quad \beta(2) = \frac{\phi(3) \cdot \mu_{00}}{\phi(2) \cdot \phi(1)}$$

$$\beta(3) = \frac{\phi(4)}{\phi(3)} \quad , \quad \beta(4) = \frac{\sqrt{\phi(5)}}{\phi(4)}$$

$$\beta(5) = \frac{\phi(6)}{\phi(4) \cdot \phi(1)} \quad \beta(6) = \frac{\phi(7)}{\phi(5)}$$

Section 3

SAR-IMAGE ANALYSIS USING REGION-BASED DESCRIPTORS

3.1 SAR IMAGE DESCRIPTION

In order to evaluate each of the three region-based change-detection methods described in Section 2, the algorithms developed were applied to synthetic-aperture radar image data furnished by RADC. Some preprocessing of all images was required before the algorithms could be applied. A brief description of the images is given here, and some of the processing steps required for analysis are described in the following sections.

The RADC supplied images consisting of 14 image files (two matching sets of 7 files each). The images correspond to data acquired by airborne radar flown at various altitudes and directional headings. The radar data was processed by a SAPPHIRE real-time digital signal processor, and is made available in the form of a sequence of digital data files on magnetic tape. The nominal ground resolution in these images is about 10 ft/pixel. Each image has 1000 lines of data, with 1026 pixels/line. Table 3-1 lists the various altitude, and headings for each image frame of the two sets of SAR images.

In our analyses, the original SAR images were averaged down to 512 x 500 pixels, with a 20 ft/pixel resolution. For 256 x 256 images this is equivalent to an area of 0.82 x 0.84, or 0.69 $(NM)^2$.

Table 3-1

SAPPHIRE SAR IMAGE DATA USED FOR CHANGE DETECTION

Serial No.	File	Flight/Pass	Heading (deg)	Altitude (ft)
1	1	1/8	272	10,200
2	2	1/9	282	7,700
3	3	1/10	272	7,700
4	4	1/11	272	7,700
5	5	1/12	270	7,700
6	6	1/13	266	7,700
7	7	1/14	262	7,700
8	8	2/1	262	7,700
9	9	2/2	266	7,700
10	10	2/3	270	7,700
11	11	2/4	272	7,700
12	12	2/5	272	7,700
13	13	2/6	282	7,700
14	14	2/7	272	10,200

Flight 1: Mission 7168 of 6/17/77
Flight 2: Mission 7181 of 6/30/77

Image pairs are (K, 15 - K), K = 1, --- 7 .

When displayed on a video monitor these images show a number of ground features. These include rows and stands of trees, a number of large buildings and other fixed structures, a set of rows of radar marker targets, and also some large areas of SAR shadow. Some of the ground features were of use in registering approximately the various image frames with one another.

In addition to these features, each image frame also contains a number of radar targets of interest for change detection. One set of targets appears in the same locations on the ground in each image frame. These targets are referred to as static targets. A second set of targets, called the movable targets, occupies a different set of locations in each of the two sets of images. Ground-truth data was made available which accurately describes the location on the ground and the nature of each target of interest.

3.2 IMAGE PREPROCESSING

The RADC supplied computer-compatible tape (CCT) was unpacked to yield 1026 x 1000 8-bit words according to the data-formatting description supplied by RADC. These pictures, when viewed directly, appear to be of limited dynamic range. In addition, several frames have data drop-outs of one or more lines where the data appeared to be all 1's.

Transforming each frame (on a pixel-by-pixel basis) according to the exponential transformation

$$y_n = 1.0331406 \text{ ** } x_n \text{ ,} \qquad (3\text{-}1)$$

as suggested by RADC, yields improved pictures in terms of contrast range and visibility of features. However, this transformation does increase the noise levels in these pictures also. This fact is of some importance since an analysis of these images subsequent to applying the change detection algorithms yielded the conclusion that the image signal-to-noise ratios were not acceptable (see Section 3.4) for an adequate test of these methods.

As a final step in the preprocessing of these images, the 8-bit results obtained from the above exponentiation operation were linearly stretched to 9 bits. This was done in order to take advantage of the 9-bit image-memory dynamic-range capability available. The images resulting from the above set of transformations were used for each of the three change-detection analysis methods.

As an example of the quality of SAR images resulting from these preprocessing operations, shown in Fig. 3-1 are a pair of images of approximately the same area of the terrain. Some of the features of interest discussed in Section 3.1 may be noted in both image frames. These photographs are representative of the quality of the available images and were used in the change-detection analysis. Figure 3-1a and 3-1b represent the arrangement of targets on two separate days, referred to as Situation 1 and Situation 2. A number of changes in target configuration are visible between the two images.

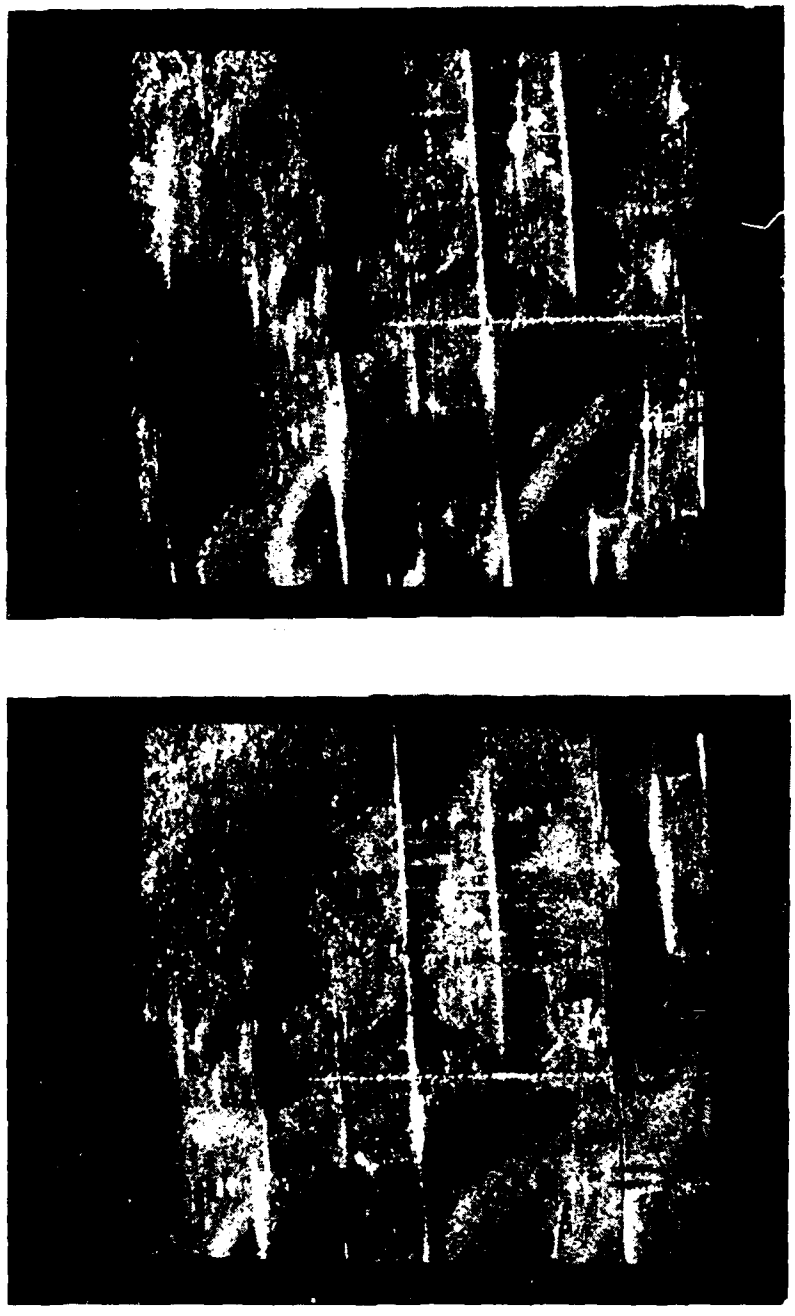


Figure 3-1 Representative SAR Imagery for Change Detection

3.3 IMAGE-QUALITY EVALUATION BASED ON AVAILABLE GROUND-TRUTH INFORMATION

In a preliminary visual inspection of the digitized image frames, some difficulty was experienced in detecting a number of the static targets. In addition, a number of the movable targets appeared in some frames, but not in others.

In order to determine possible reasons for the inability to detect some of the targets, an analysis of the images was performed using available ground-truth information about the targets supplied by RADC. This information was quite detailed, and included a specific description and photographic views of each target, and a location specification (in terms of latitude and longitude) which enabled each target to be located to within a pixel of its position on the ground. Targets were also specified as being either static or movable targets. Also available was a copy of the analog, SAR-imagery film containing the original image data.

Inspection of the original analog photographic data showed that the digitized images were not significantly worse in image quality than the originals. The remainder of the analysis was therefore confined to the digitized images provided. Of the 14 images supplied, 7 were selected for analysis. For each of these, the estimated locations of all appropriate static and movable targets were computed and plotted. Estimates of the pixel locations of all targets were made by a least-squares method of analysis. In this procedure, as many strong targets as possible are first identified in a particular image frame. Each target's latitude and longitude are determined from ground-truth information, and the pixel location of each is measured on the video display by a pair of cursors. Using the list of pixel locations and latitude/longitude values, a solution is next obtained to the coefficients of the pair of equations.

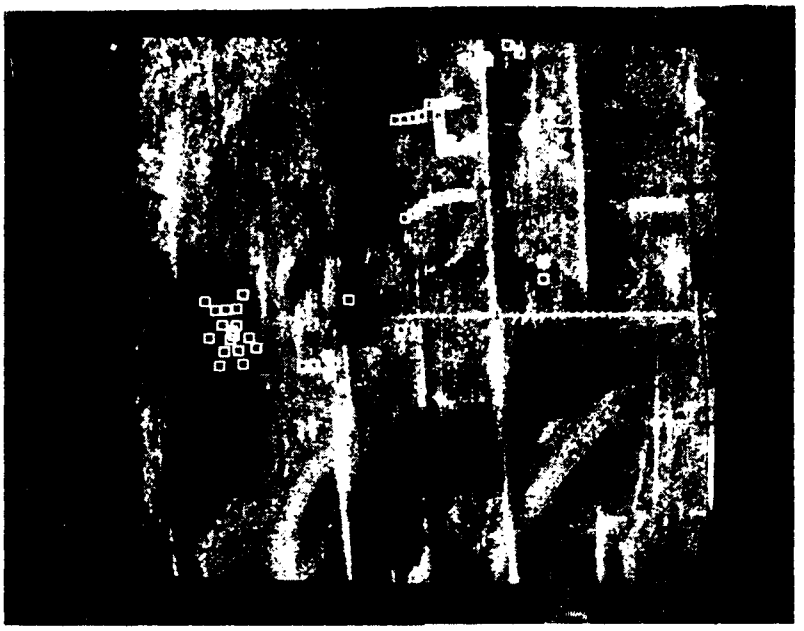
$$u = ax + by + c \tag{3-2}$$

$$v = dx + ey + f \tag{3-3}$$

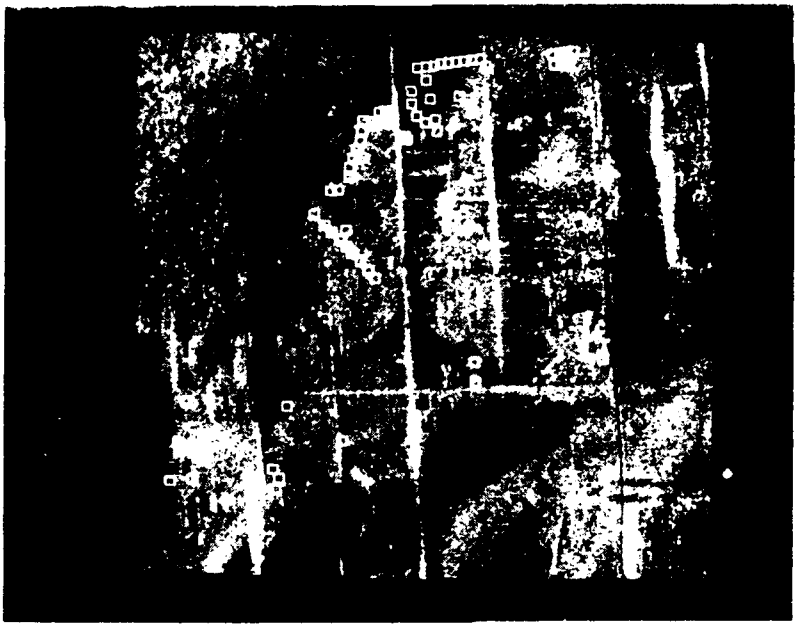
where the coefficients a, b, c, d, e, and f are computed so as to minimize the value of an error term. If more than three such targets are used, the solution to the above equations is overdetermined, and both a best-fit estimate of the coefficient values, as well as an estimate of the error for the fit are obtained from this procedure. Between five and fifteen targets were used in the least-squares fitting procedure for each image frame, and in each case the residual r.m.s. error was less than one pixel in each direction. These results seem to show that the constraints required on the altitude and heading during data acquisition had been met very well.

The computed coefficients used to transform latitude and longitude to pixel location in the image frame were then used to plot the estimated locations of all targets in each image frame. As an example, Figure 3-2 shows the estimated locations of all targets within each image frame for the two (Situation 1 and Situation 2) images shown in Figure 3-1.

This procedure revealed that the primary reason for the inability to detect a number of the static targets is that they are located in areas of SAR shadow in the furnished images. In addition, most of the targets of interest are not of sufficient



a



b

Figure 3-2 Estimated Locations of Targets in Situation 1 and Situation 2 Images

extent or strength to be distinguished clearly from the background. A reason for this may be the fact that the static targets are, in general, smaller in size than the movable targets in both image sets.

Table 3-2 summarizes the visibility of static targets in a representative image frame. Ground-truth data indicates that targets numbered 20 through 27 in the table are all small vehicles.

The procedure described previously for determining the visibility of targets was performed on the pixel-averaged versions of an image frame, so that the entire image could be displayed as in Figures 3-1 and 3-2. However, since many of the targets appear to be difficult to see in these averaged-down images, the above procedure was also applied to a segment of the full-resolution SAR image. For this purpose, a segment known to contain the least visible set of fixed targets was selected.

Figure 3-3 shows a 512 x 512 pixel section of the original full-resolution, SAR image. Many of the brighter objects, which are buildings and other structures, can be identified with objects in Figures 3-1 and 3-2. The particular group of targets of concern was the set of fixed targets numbered 20 through 27 in Table 3-2. These targets are clustered together in the general area of the SAR shadow which can be seen at the top center of Figure 3-3. In this area one strong target that can be seen, is identifiable with target No. 18; another target can be seen, but is very weak in comparison with the returns from the movable targets that can be seen in various parts of the picture. Targets 20 through 27 as a group are not visible in this full-resolution image, just as they are invisible in the reduced resolution images of Figures 3-1. Therefore, we conclude that the problem in seeing some of the targets of interest did not arise because of the averaging process, but rather was due to the quality of the original SAR imagery.

It is of interest to note that a pronounced fringe pattern is visible over the entire image frame in each of the full-resolution, SAR images. This characteristic makes each target, both strong and weak, appear to be a multilobed structure, with multiple sidelobes in the range direction.

3.4 EVALUATION OF IMAGE QUALITY BY STATISTICAL ANALYSIS

Because of the poor visibility of many of the targets in the furnished SAR imagery, some further evaluation of the quality of these images was considered desirable. As a further check, some experiments were devised for this purpose, and some statistical analysis of the image data was performed. These experiments are described below.

When the set of images provided has been subjected to identical intensity transformations, and have been registered with one another, several of these images represent a picture of the same piece of terrain. In a noise-free situation they should be identical; this is not the case for real images. However, images from the same set should be sufficiently identical to permit two image segments to be correlated, and a high-correlation coefficient to be obtained.

The situation may be described analytically as follows. Given two identical segments of two image frames of the same ground area, the set of pixels in the two segments can be described by $(S+N_1)$ and $(S+N_2)$, where the component S is the

Table 3-2

SUMMARY OF VISUAL-STATIC TARGET DETECTIONS
IN PIXEL-AVERAGED SAR-IMAGE FRAME

Static Target No.	Can Be Seen	With Difficulty	Cannot Be Seen	Comments
1			x	In SAR Shadow
2		x		
3		x		
4	x			
5		x		Edge of SAR Shadow
6		x		Edge of SAR Shadow
7		x		Edge of SAR Shadow
8		x		Not in SAR Shadow
9		x		Not in SAR Shadow
10		x		Not in SAR Shadow
11		x		Not in SAR Shadow
12			x	Not in SAR Shadow
13			x	In SAR Shadow
14		x		
15	x			
16	x			
17	x			
18	x			Edge of SAR Shadow
19			x	Obscured by Stronger Reflections From Marker Targets
20			x	
21			x	In SAR Shadow Region
22			x	
23			x	
24			x	
25			x	Edge of SAR Shadow
26			x	
27			x	In SAR Shadow



Figure 3-3 Full-Resolution 512 x 512 Pixel Segment of SAR Imagery for Target-Visibility Experiment

"signal" common to both images, and N_1 and N_2 are the "noise" components. It is assumed that the noise in each image is independent and stationary. It is required to form an estimate of the signal-to-noise ratio S/N_1 (or S/N_2) in each image.

For analysis, we compute the sum and difference images by adding (subtracting), on a pixel-by-pixel basis, the two image segments. We also compute the means and standard deviations of these sum and difference images, as well as those of the two image segments themselves. For this purpose, the set of pixels in each image segment is treated as a one-dimensional array.

The variance in each image segment (1 and 2) can be described as:

$$\sigma_1^2 = \sigma_s^2 + \sigma_{n_1}^2 \quad , \tag{3-4}$$

and

$$\sigma_2^2 = \sigma_s^2 + \sigma_{n_2}^2 \quad . \tag{3-5}$$

When the sum and difference images are formed, their variances are given respectively by:

$$\sigma_{sum}^2 = 4\sigma_s^2 + \sigma_{n_1}^2 + \sigma_{n_2}^2 \quad , \tag{3-6}$$

and

$$\sigma_{diff}^2 = \sigma_{n_1}^2 + \sigma_{n_2}^2 , \tag{3-7}$$

due to the fact that the signal component adds directly in the sum image whereas the noise component is independent in the two images. Now, assuming that $\sigma_{n_1}^2 = \sigma_{n_2}^2 = \sigma_n^2$, we get:

$$\frac{\sigma_{sum}^2}{\sigma_{diff}^2} = \frac{2\sigma_s^2 + \sigma_n^2}{\sigma_n^2} \tag{3-8}$$

or, that

$$\frac{\sigma_s}{\sigma_n} = \left[\frac{\sigma_{sum}^2 - \sigma_{diff}^2}{2\sigma_{diff}^2} \right]^{1/2} \tag{3-9}$$

Which is the required expression for the signal/noise ratio in these images.

An alternative way of estimating the signal-to-noise ratio is based on calculating the correlation coefficient ρ between the two image segments, again considering each image to be a one-dimensional array of pixel values. The correlation coefficient ρ can be related to the scene variance as follows.

$$\sigma_{11}^2 = \sigma_s^2 + \sigma_{n_1}^2 \ , \quad \sigma_{22}^2 = \sigma_s^2 + \sigma_{n_2}^2 \ , \tag{3-10}$$

and

$$\sigma_{12}^2 = \sigma_s^2 \ ,$$

where

σ_{11} and σ_{22} refer to image 1 and 2.

From these,

$$\rho = \frac{\sigma_s^2}{\left(\sigma_s^2 + \sigma_{n_1}^2\right)^{1/2} \left(\sigma_s^2 + \sigma_{n_2}^2\right)^{1/2}} \ , \tag{3-11}$$

and if we assume $\sigma_{n_1}^2 = \sigma_{n_2}^2$, then

$$\rho = \frac{\sigma_s^2}{\sigma_s^2 + \sigma_n^2} \tag{3-12}$$

Therefore,

$$\frac{\sigma_s}{\sigma_n} = \sqrt{\frac{\rho}{1 - \rho}} \ . \tag{3-13}$$

Equating the right-hand sides of Equations (3-13) and (3-9), it is possible to obtain

$$\rho = \frac{\sigma_{sum}^2 - \sigma_{diff}^2}{\sigma_{sum}^2 + \sigma_{diff}^2} \ , \quad \text{and} \quad \frac{\sigma_{sum}^2}{\sigma_{diff}^2} = \frac{1 + \rho}{1 - \rho} \ . \tag{3-14}$$

Thus, the image quality, in terms of the above defined signal-to-noise ratio, can be calculated from either the sum and difference image variances, or from the correlation coefficient.

As an example of these calculations, the variances and correlation coefficient were calculated for a pair of representative SAR images. A 128 x 128 segment, chosen to contain no movable targets, was selected from each image and sum and difference variances were computed.

Table 3-3 summarizes the results of these calculations, and shows the obtained estimates of the image signal/noise ratio. These numbers are typical of the set of SAR images supplied. It is seen that the signal/noise ratio is estimated to be only about 0.7.

Table 3-3

SIGNAL-TO-NOISE RATIO ESTIMATES
FOR FULL-RESOLUTION, SAR IMAGES

Full-Resolution Image Size: SAR Image 128 x 128 pixels

	Mean Value	Standard Deviation
Image 1	147.88	31.22
Image 2	146.80	31.16
Sum Image	294.67	50.54
Difference Image	1.083	36.57

Correlation ρ between images: 0.3126

Calculated Signal/Noise Ratio

Based on Image Variances: 0.674 (Eq. 3-9)

Based on Correlation Value: 0.674 (Eq. 3-13)

From the foregoing calculations, we conclude that the signal-to-noise ratio in the furnished SAR images is very poor, and that these images are inadequate for a conclusive test of the area-based change-detection algorithms.

In order to improve the signal-to-noise ratio, a method of signal averaging was tried on each of the image frames. In this method a new image is formed by averaging a set of 2 x 2 adjacent pixels in the original image to form one pixel in the new image. The resulting picture is half the size of the original picture in each dimension (that is, contains one quarter of the original number of pixels). The averaged images for two of the supplied image frames were computed, and the signal-to-noise ratio analysis described previously was applied. The results of this experiment are summarized in Table 3-4.

The results in Table 3-3 and 3-4 show that the signal-to-noise ratio in the furnished, SAR images is less than unity. The situation can be improved by a method of pixel averaging, but the resulting image is still very noisy. A problem with the pixel averaging process is that targets of interest which were a few pixels in extent are now only one or two pizels in size in the averaged images. Thus an improvement in the signal-to-noise ratio has been obtained by sacrificing resolution, which is now approximately 20 ft/pixel in the reduced images. For this reason, area-based, change-detection schemes will not necessarily work better with pixel-averaged images (see Section 4.3).

Table 3-4

SIGNAL-TO-NOISE RATIO ESTIMATES
FOR REDUCED-RESOLUTION SAR IMAGES

Reduced-Resolution SAR Image
Image Size: 64 x 64 pixel

	Mean Value	Standard Deviation
Image 1	147.51	24.26
Image 2	146.42	24.00
Sum Image	293.93	42.05
Difference Image	1.09	23.67

Correlation ρ between images: 0.5188

Calculated Signal/Noise Ratio

Based on Image Variances: 1.04 (Eq. 3-9)

Based on Correlation Value: 1.04 (Eq. 3-13)

As a final experiment designed to show conclusively the importance of the image signal-to-noise ratio in these change-detection schemes, two, identical SAR images were compared for change after independent noise of known mean and variance was added to each image segment

Two, identical image segments of the full-resolution SAR image were used for this purpose. (These image segments are the same as that referred to as Image 1 in Table 3-3.) It was found that independently added noise with a standard deviation of 60.0 was needed to degrade the signal-to-noise ratio to a point where the correlation between the image segments was as poor as that for two frames of the furnished SAR imagery.

The results of adding this controlled amount of noise, and then performing the signal-to-noise ratio calculations described before, are shown in Table 3-5. It should be pointed out that the noise added in this experiment was wideband, that is, the noise

samples were independent from pixel to pixel. This is not quite the kind of "noise" that contributed to differences between nominally identical frames in the supplied SAR images, where in addition to some scene noise, real differences in scene content (contributing to noise at low frequencies) gave rise to poor correlation between images frames, and thus degraded the signal-to-noise ratio.

Table 3-5

SIGNAL/NOISE RATIO ESTIMATES FOR TWO IDENTICAL SAR IMAGE SEGMENTS TO WHICH A CONTROLLED AMOUNT OF NOISE HAS BEEN ADDED

Full-Resolution SAR Image
Image Size: 128 x 128 Pixels

	Mean Value	Standard Deviation
Image 1	253.9	62.53
Image 2	253.2	62.61
Sum Image	507.1	99.57
Difference Image	0.727	75.82

Correlation ρ between images: 0.266

Calculated Signal/Noise Ratio

Based on Image Variances: 0.60 (Eq. 3-9)

Based on Correlation Value: 0.60 (Eq. 3-13)

The result of applying one of the change detection methods – the phase-correlation algorithm – to the noise-added image pair indicates that, while most of the elementary regions (32 x 32 pixels in size) are identified by the algorithms as being approximately registered, the average peak amplitude obtained for the primary peak is only about 0.27. This is in accordance with the known signal-to-noise ratio in the images. However, in some cases, the algorithm indicates that a displacement of several pixels exists between corresponding elementary regions. Thus, we are at the threshold of good performance for the phase-correlation algorithm when the signal-to-noise ratio is degraded to 0.6. Since this is comparable with the actual signal-to-noise ratios in the furnished SAR images, we conclude that the available image quality falls short of being acceptable for verification of the change-detection schemes. It may be noted that the results from the above experiment, where wideband noise is added to the image, will be better than the case in which an equivalent amplitude image noise resides primarily in the distortions and variations in scene content.

3.5 IMAGE SEGMENTATION SCHEME

As described in Section 1.2, the use of each of the three methods of region-based, change detection involves first defining a set of regions in each image. These regions are a contiguous set of pixels and occupy the same location in each image frame to be

compared. A number of schemes may be used to define the set of regions used for change detection; the particular scheme used in the present study is described here.

Figure 3-4 shows a small section of two images being compared. For convenience, each image is divided into a number of square areas adjacent to one another. Also, for convenience in computing Fourier transforms, each image as well as each of these squares was chosen to be of a size which is an integral power of 2 on a side. Typically, each image is 256 x 256 pixels, while each elementary square is 16 x 16 or 32 x 32 in size.

The elementary regions on which the various descriptors were calculated consist of these equare areas which are 16 x 16 or 32 x 32 pixels in size. However, adjacent regions are not the adjacent squares, as shown in the figure. Rather, the adjacent elementary region is that which overlaps the previous region by half the number of pixels per region. This scheme was used, so that pixels which are on the border of a given region would be in the interior of some other adjacent region. Similarly, the next row of elementary regions also overlaps regions in the previous row by half the number of pixels.

The number of elementary regions for which the similarity measures for each of the three change-detection methods need to be calculated depends on the size of each elementary region, as well as on the size of each image. As an example, for an image size of 256 x 256 pixels, and an elementary region of size 32 x 32 pizels, there will be 15 x 15 , or 225 elementary regions defined. In general, for an image size of $2^n \times 2^n$ pixels, and a region size of $2^k \times 2^k$ pixels, the number of regions is $\left[2^{n+1}/2^k - 1\right]$.

The following considerations govern the choice of size of each elementary region. As the size of each region is made smaller, the detection of possible changes is indicated with greater resolution, since the number of elementary regions is increased. At the same time, a larger number of descriptors (one set for each region) must now be be calculated, and, thus, there is an undesirable decrease in the capacity for data reduction which region-based descriptors can provide. A second consideration determining cell size is a statistical one. In the presence of image noise, it is desirable to have regions that are large, in order that the computed values of the descriptors used should be noise-averaged over as many pixels as possible.

A third consideration concerns the possibility of imperfect overlap between two, elementary regions to be compared, because of approximate initial registration, or some rotation between two images which are otherwise well-registered over some small area. The size of an elementary region in this case needs to be finite even though the image signal-to-noise ratio is very high, because of possible edge effects as seen by an "invariants" descriptor.

As in the case of image noise, misregistration can also be described in terms of an effective "signal-to-noise" ratio. If an N x N scene is shifted in each direction by ℓ pixels, the common area between the two segments is now:

$$\omega = (1 - \ell/N)^2 \quad . \tag{3-15}$$

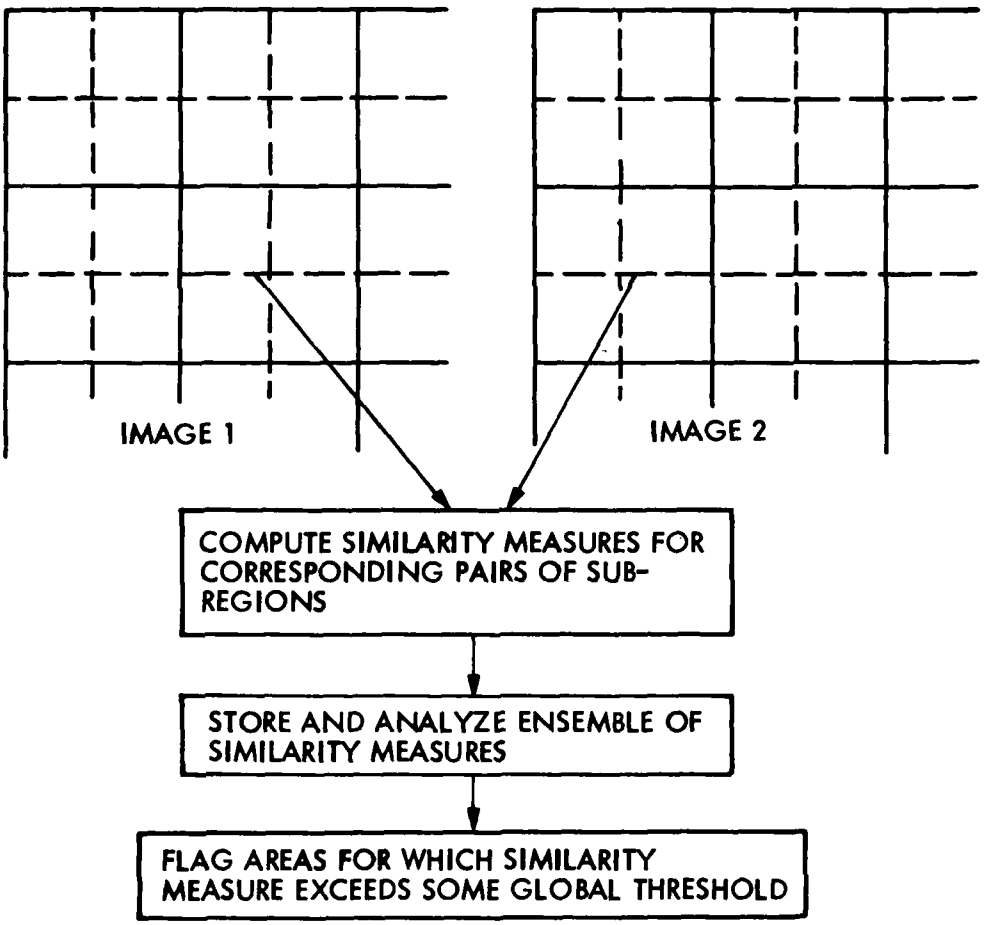


Figure 3-4 Image-Segmentation Scheme and Processing Sequence

Assuming a stationary scene, the "signal" power is proportional to ω, while the "noise" power is proportional to $(1-\omega)$. Thus the effective signal-to-noise ratio as seen by an invariants descriptor will be

$$SNR = \left(\frac{\omega}{1-\omega} \right)^{1/2} = \left[\frac{\left(1 - \ell/N^2\right)}{1 - \left(1 - \ell/N^2\right)} \right]^{1/2}$$

$$= \left[\frac{r^2}{1-r^2} \right]^{1/2}, \text{ where } r. = (1 - \ell/N) \qquad (3\text{-}16)$$

is the fractional overlap ratio in each direction.

The manner in which this "signal-to-noise" ratio behaves as a function of fractional-overlap ratio is shown in Figure 3-5. As an example, for a shift of 2 pixels in a 32 x 32 image segment, the signal/noise ratio is approximately 2.7; for a 2 pixel shift in a 16 x 16 image, the S/N ratio is only 1.8. For the signal/noise ratio to be better than, say, 3, the corresponding registration must be achieved such that the fractional overlap is better than 0.9 in each dimension. These results show that, if the possible lack of registration of 1 or 2 pixels is not to be a limiting factor in the comparison of two, image segments, then the size of each segment must be about 32 x 32 pixels or larger.

A good choice for the size of regions to be used can be made if the sizes of the targets of interest are known beforehand. In that case, the elementary regions can be chosen to be of a size comparable to the size of these targets, or somewhat larger, unless considerations of noise or overlap prevail. In the present case, all targets of interest were comprised of 5 (or less) pixels; however, the conclusions from the previous discussion of the overlap problem, and the presence of significant image noise (Section 3.4) dictated the use of regions no smaller than 32 x 32 pixels.

3.6 DATA PROCESSING SEQUENCE

The preliminary processing steps and the segmentation scheme for defining elementary regions were discussed in the previous subsections. The sequence of operations required for the automatic detection of change between a pair of images can now be summarized:

a) Align images to within the required tolerance

b) Select a size for the elementary regions to be processed

c) For each elementary region, compute descriptors for each of the three algorithms,
   - phase correlation
   - power-spectral coefficients
   - moment invariants

d) For each algorithm, compute a similarity measure for each elementary region

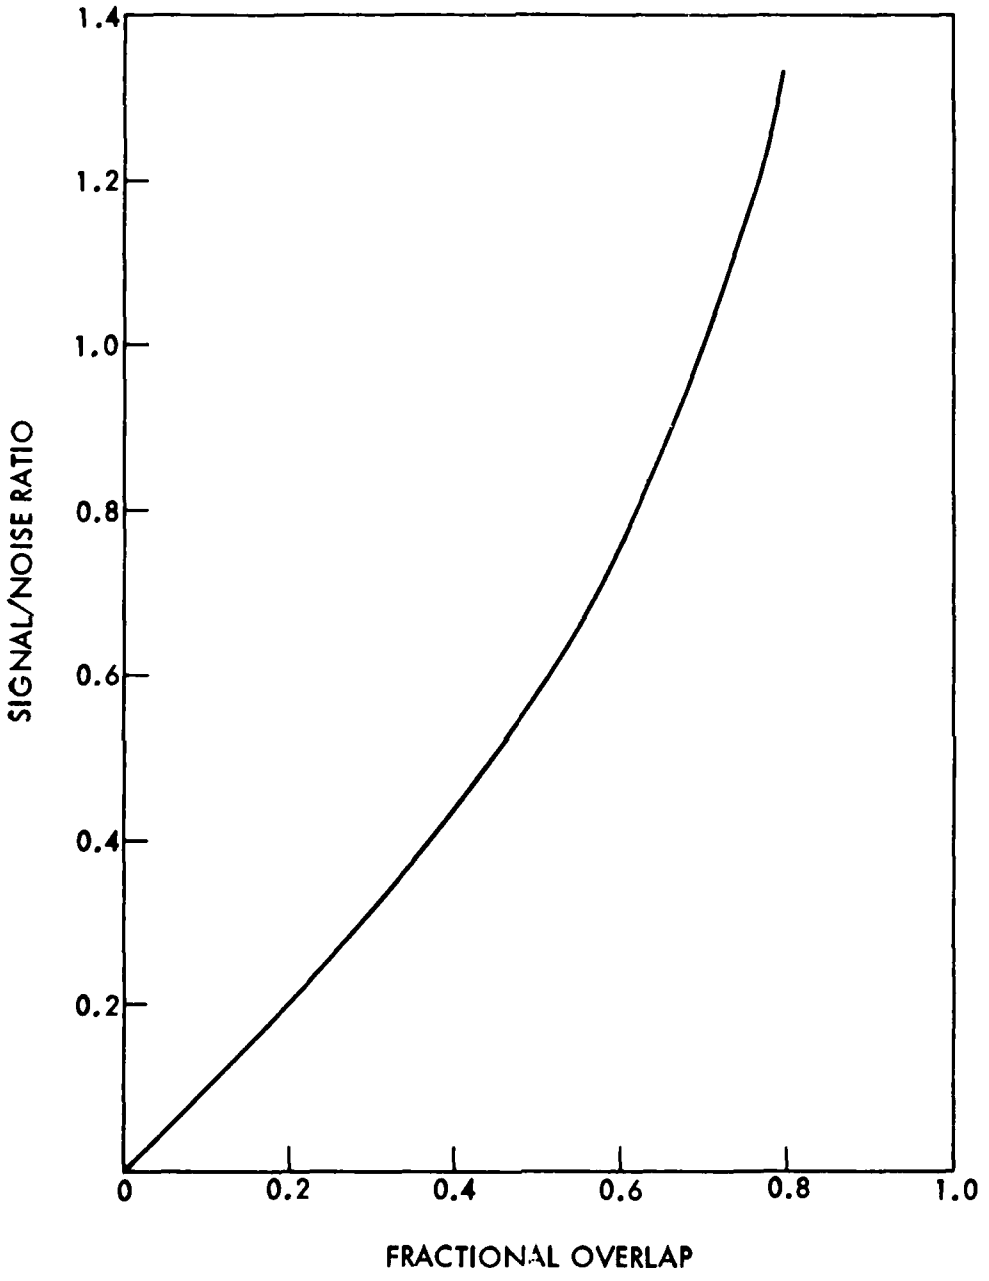


Figure 3-5 Effective Signal/Noise Ratio as a Function of Lack of Overlap Between Two-Image Segments for Invariant Descriptors

e) Examine the global statistics for each set of similarity measures and determine global thresholds

f) Flag those areas for which the similarity measure fails to exceed the global threshold, as being likely areas of change.

In this section, the processing details for each of the three algorithms are discussed.

Images were aligned prior to change detection by a least-squares fitting procedure similar to that described in Section 3.3. A number of prominent features common to both images of a pair to be compared were selected, and coefficients for the best fit to align one image with the other were calculated. Because of possible problems with any interpolation schemes for resampling, the comparing of two images was confined to those having the same nominal directional heading. In this way, images could be aligned by specifying only a vector translation between image frames. After alignment, sections of each image from identical locations were selected for further analysis.

The tolerance required for the initial approximate alignment of images is a function both of image detail and image noise. Because of the poor signal/noise ratio in these images, the question of alignment tolerance was analyzed in some detail. These details are described in Section 4.1. The selection of size of an elementary region for processing is also dependent of the image noise, and has been discussed in Section 3.5.

Also, because of the signal/noise problem with these images (and because of the availability of a 512 x 512-pixel image memory), the change-detection algorithms were applied to images which were pixel-averaged versions of the original image. This averaging procedure has been described in Section 3.4. This has the advantages of somewhat improving the signal-to-noise ratio; also a larger segment of each image could be processed at one time. A disadvantage now is the reduced resolution of targets of interest.

A 256 x 256-pixel area of each image to be compared was selected. The elementary region for each of the algorithms was 32 x 32 pixels. Using this scheme, a total of 15 x 15, or 225 similarity measures are thus computed over the area of each image pair analyzed.

For phase correlation, each pair of elementary image segments to be compared was first Fourier transformed. A cosine-weighting window was applied to the data for this purpose. The cosine weighting minimizes the contribution of the dissimilar parts of the two segments at their edges, especially when there is some lack of overlap. Cosine weighting is also useful in reducing the effect of any image shading which may be present. An effect of the weighting is that the primary peak is somewhat broadened in the phase-correlation output. This fact is of minor importance, since the peak location is determined with sub-pixel accuracy by a process of quadratic interpolation in the developed phase-correlation algorithms used here. For each calculated primary peak amplitude and vector displacement, a similarity measure based on the definition in Equation 2-3 was calculated, using a σ value of 1. (The displacement here is expressed in pixels, and p is a normalized, peak height with a

value in the range 0 to 1.) Based on the image-segmentation scheme used, these calculations gave a matrix of values of the similarity measure for phase correlation, a matrix which we term the phase-correlation change matrix.

For power-spectrum change detection, each elementary region to be compared was again first cosine weighted, then Fourier transformed to obtain power-spectral values of a matrix of 32 x 32 values, all elements outside of a 21 x 21 region around the origin were discarded. Also, a region of size 9 x 9 pixels around the origin was set to zero. This procedure leaves a set of 180 significant power-spectrum values for further analysis. Thus, a reduction of the data by a factor of approximately 0.18 is achieved in this process.

Similarity measures for each region were computed by calculating the correlation coefficient between the set of power-spectral values for each image. This procedure has the advantage that the image regions are not required to be normalized or contrast-equalized prior to doing the analysis. This procedure also yields a similarity measure which is in the range 0 to 1, and, again, a matrix of 225 values were obtained for the entire image. We term this matrix the power-correlation change matrix.

For the third method using invariant moments, cosine weighting was used also. This has the effect of minimizing the effect of strong features which appear at the edges of an elementary region, since in moment computations, the edges are strongly weighted over the remaining parts of the region. A series of fast-vector product operations were used to generate the various moments, and the six invariants as described in Table 2-2 were calculated.

The similarity measure in this case consists of the correlation coefficient between the six invariants of the corresponding elementary regions in each image. Again, a number normalized to the range 0 to 1 is obtained, and a matrix of 225 similarity measures is generated. This matrix is called the invariants-correlation change matrix.

3.7 DATA ANALYSIS

The results of applying each of the three change-detection algorithms to a pair of corresponding images as described in the previous section is a set of three "change matrices", each of which has 225 elements. In an operational change-detection processor, these matrices would be subjected to some thresholding operation (the threshold having been previously determined, or being calculated based on the statistics of the ensemble of similarity measures, as described in Section 3.6).

In the present study, it was desirable to determine some method by which the results of the three change-detection schemes could somehow be related to our knowledge of the location of targets on the ground, in order to obtain an absolute measure of how well each algorithm performed. Therefore, instead of an analysis of each change-detection matrix separately, a method for comparing each change matrix with an "ideal" matrix of known target changes was devised. Also, each change matrix could also be correlated with each of the other two, in order to determine whether the three methods were in some sense consistent with one another in their performance.

The ideal or theoretical "change matrix" is generated from a knowledge of the pixel locations of all movable targets (this is known from available ground-truth data) in the following ways: an image-segmentation scheme identical to that used for each change-detection scheme is first defined. Each movable target is then mapped into one or more of the defined, elementary regions. A map of the number of targets in each region is thus created, and such a map is generated for each of the two Situations, 1 and 2, for which images are available.

Each map of the number of targets is also a 15 x 15 matrix, in which elementary regions with no targets in them have a value of zero; other regions have a value equal to the number of targets present. We now take the absolute difference between the target map for Situation 1 and the target map for Situation 2. (Because of this differencing, only the movable targets need be considered in defining the theoretical-change matrix.) The resulting 15 x 15 matrix of values is in some sense the ideal pattern of change against which any derived change detection results could be compared. We call this matrix the theoretical change matrix.

Analysis of the data now consists of correlating each of the three derived change matrices with the above theoretical change matrix. An example of the theoretical change matrix for a 256 x 256 image segment, with a cell size of 32 x 32 pixels, is shown in Figure 3-6. Also, in Figure 3-7 is shown the change-detection matrix derived from applying the phase correlation algorithm to the same image areas.

Finally, the results obtained from correlating each of the three change-detection matrices against the theoretical-change matrix, as well as against one another, are summarized in Table 3-6. We note that the three methods seem to be consistent in their results, but have a weak correlation with actual ground-truth data. This may be due to the poor visibility of a number of the targets, as previously described (Section 3.3).

3.8 COMBINATION OF ALGORITHMS

The results obtained separately from the three, change-detection algorithms were combined, in order to determine whether such a combination could result in a more effective change-detection procedure. The results of such an experiment are described below.

For each of the three change matrices, a threshold was chosen such that a fixed number (about 35) of points exceeded the threshold. Each change matrix was then converted into a binary matrix by using this threshold value. The three binary matrices were then added to produce an overall change matrix. This resulting matrix is not a binary matrix.

The sum of the binary matrices was then correlated with each of the three individual change matrices and with the theoretical matrix of target change. The results were negative; the correlation with the theoretical change matrix is only 0.17. The correlation with each of the three individual change matrices is lower than this value. Also, the correlation between the binary versions of each change matrix was also found to be insignificant. Therefore, we conclude that we have not obtained significant results in each of the three change-detection methods applied separately, and, thus, that the method for combination of the results needs further verification.

2	0	3	1	0	0	0	0	0	0	0	0	3	3	0
2	4	5	1	0	0	0	0	0	0	0	0	1	1	0
2	4	5	2	1	0	0	0	0	0	0	0	0	0	0
4	1	6	5	2	0	0	0	0	0	0	0	0	0	0
5	3	2	3	1	0	0	0	0	0	0	0	0	0	0
1	3	3	1	2	0	0	0	0	0	0	0	0	0	0
1	2	1	2	2	0	0	0	0	3	0	0	1	4	5
3	5	3	1	0	0	0	0	0	0	0	0	1	4	5
2	2	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 3-6 Theoretical Matrix of Target Change; SAR Images from Situations 1 and 2

17	14	262	318	50	62	6	8	6	139	55	16	22	9	112
139	5	36	124	34	27	146	62	101	6	56	204	70	16	1
144	61	365	163	112	111	57	41	27	43	28	5	9	296	109
4	56	121	120	14	20	7	10	64	144	120	104	28	192	20
11	2	441	27	2	339	26	26	39	171	32	210	30	26	15
213	3	36	19	4	9	9	60	2	71	127	234	39	79	153
11	127	211	197	6	25	17	153	167	2	171	7	19	6	16
216	62	84	25	61	236	206	80	89	5	121	17	24	1	146
20	8	90	159	184	36	270	65	337	123	145	31	36	79	3
255	225	91	150	5	17	64	10	129	227	43	31	20	99	14
16	303	207	1	47	279	16	328	256	290	253	28	510	317	18
236	259	3	74	410	24	3	9	30	24	313	1	377	399	6
152	349	169	366	109	302	283	292	331	154	22	209	216	312	113
317	233	421	245	324	451	376	331	251	451	222	479	87	201	219
346	132	16	391	438	45	11	14	362	52	251	163	71	7	357

Figure 3-7 Phase-Correlation Change Detection Matrix
SAR Images From Situations 1 and 2

Table 3-6

RESULTS OF CORRELATION BETWEEN THE VARIOUS CHANGE-DETECTION TECHNIQUES

	PHASE CORRELATION CHANGE MATRIX	POWER CORRELATION CHANGE MATRIX	INVARIANTS CORRELATION CHANGE MATRIX
THEORETICAL CHANGE MATRIX	0.4843	0.5028	0.6552
PHASE CORRELATION CHANGE MATRIX		0.9577	0.9083
POWER CORRELATION CHANGE MATRIX			0.9376

3.9 DETECTION RATE AND FALSE-ALARM ESTIMATES

It must be emphasized that the three algorithms which were discussed are all algorithms for change detection, and not for the detection of targets as such, in imagery. These approaches are intended solely to give a reliable indication either of the presence of a previously absent target, or the absence of one previously present. If a strong target is present in the same location of two image frames to be processed, each of these techniques will (quite correctly) not indicate presence of such a target.

In view of this fact, questions regarding detection rate and false-alarm probability must be interpreted somewhat differently from the conventional situation when targets are being detected. In our approach the emphasis is on detecting regions of change (either correctly or incorrectly), and thus any estimation of detection probability should include this consideration. (By prior discussion with RADC, region-based criteria for determining these error rates were determined as being more appropriate to our investigation than conventional measures.)

As an example, the detection rate for regions of change is computed as follows. Figures 3-8, 3 9, and 3-10 show three maps which are the outputs from each of the three change-detection algorithms. These figures represent regions of the image area which were analyzed, and those regions labeled 1 are areas in which significant change was determined to have occurred. Figure 3-8 represents the output of the phase correlation algorithm, with 34 detected instances of change; Figure 3-9 is the output from the power spectrum correlation method, from which 33 changes were found; and Figure 3-10 shows 34 areas of change as determined by the invariant moments approach. For reference, compare the target difference map as generated from available ground truth data (Figure 3-6).

With reference to Figure 3-6, the fraction of number of correct areas of change detected by these algorithms can be computed as 22% for phase correlation, 28% for power spectrum correlation, and 14% for the invariant moments correlation, respectively. Because the outputs from these algorithms are based on comparisons of identical regions on the ground rather than of specific targets, we believe that the answer to questions such as "number of targets detected" can be given only on a statistical basis, as above.

The number of false alarms (again referred to the fraction of regions incorrectly labeled rather than to the number of false targets) is also estimated in a similar way. Referring to the Figures 3-8, 3-9, and 3-10, regions which are mislabeled in these as areas of change are included in the false-alarm figure; the rate turns out to be 23 false alarms for the phase correlation method, 19 for the power spectrum method, and 27 for the invariant moments method. Since the region over the total image is of area 0.69 $(NM)^2$ (Section 3-1), this amounts to false-alarm rates of $25.2/(NM)^2$, $27.7/(NM)^2$, and $39.2/(NM)^2$, respectively, for these three techniques.

0	0	0	0	0	0	1	1	1	0	0	0	0	1	0
0	1	0	0	0	0	0	0	0	1	0	0	0	0	1
0	0	0	0	0	0	0	0	0	0	0	1	1	0	0
1	0	0	0	0	0	1	0	0	0	0	0	0	0	0
0	1	0	0	1	0	0	0	0	0	0	0	0	0	0
0	1	0	0	1	1	1	0	1	0	0	0	0	0	0
0	0	0	0	1	0	0	0	0	1	0	1	0	1	0
0	0	0	0	0	0	0	0	0	1	0	0	0	1	0
0	1	0	0	0	0	0	0	0	0	0	0	0	0	1
0	0	0	0	1	0	0	0	0	0	0	0	0	0	0
0	0	0	1	0	0	0	0	0	0	0	0	0	0	0
0	0	1	0	0	0	1	1	0	0	0	1	0	0	1
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	1	0

Figure 3-8 Regions of Change Labeled by Phase Correlation Algorithm

0	1	1	0	0	0	0	0	0	0	0	1	1	0	0
0	0	0	1	0	0	0	0	0	0	0	0	1	1	0
0	0	1	1	0	0	0	0	0	0	0	1	1	0	1
0	0	0	0	0	0	0	0	0	0	1	0	0	0	0
0	0	0	0	0	0	1	0	0	0	0	0	1	1	1
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	1	0	0	0	0	0	1	0	0	1	0	0
0	0	1	1	0	0	0	0	0	0	0	0	1	0	1
0	0	0	1	0	0	0	0	0	0	0	1	0	0	0
0	0	0	0	0	1	0	0	0	0	0	0	1	0	0
0	0	0	0	0	0	1	0	0	0	0	1	0	0	0
0	0	0	0	0	0	0	0	0	1	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	1	0

Figure 3-9 Regions of Change Labeled by Power Spectrum Correlation Method

0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
0	0	1	1	0	0	0	0	0	1	0	1	0	1	0
0	0	0	0	0	0	0	1	0	0	0	0	0	0	1
0	0	0	0	0	0	0	0	0	0	0	0	0	1	0
0	1	0	1	1	0	0	1	1	0	0	0	0	0	0
0	0	0	0	0	0	0	0	1	0	0	0	0	1	0
0	0	0	0	0	0	0	0	1	0	0	1	0	0	0
0	1	0	0	0	0	0	1	0	0	0	0	0	1	0
1	0	0	0	0	0	0	0	0	0	0	1	1	0	0
0	0	0	0	1	0	0	0	0	0	0	0	0	0	0
0	0	0	0	1	0	0	0	0	0	0	0	0	1	0
1	0	0	0	1	0	0	0	0	0	0	0	1	0	0
0	0	0	1	0	0	0	0	0	0	1	0	0	0	0

Figure 3-10 Regions of Change Labeled by Invariant Moments Method

Section 4

SENSITIVITY ANALYSIS

4.1 IMAGE OVERLAP

The previous section has described in detail the three change-detection algorithms investigated under the present study, and some of the results obtained. Although the results from all of these methods appear to be consistent, they are only weakly correlated with actual ground-truth data.

An analysis of the image signal-to-noise ratio indicated that the input images for analysis were excessively noisy (Section 3.4); also, large areas of SAR shadow caused many targets to be invisible (Section 3.3) in these images.

Given the signal-to-noise ratio available in the furnished SAR images, it is possible to analyze the images, so as to obtain an estimate of the degree of change which can be reliably detected. From the results of Section 3.4, we know that the correlation between two nominally identical areas of two SAR images is only about 0.7. We now determine the way in which the correlation of the change-detection matrix between two identical segments of the same SAR image drops off, as the overlap between these image segments is gradually decreased. This can be done for each of the three algorithms by calculating the change-detection matrix for an image segment. Then calculating the same matrix for an image segment which is displaced by one or more pixels from the first. By progressively decreasing the degree of overlap, the correlation should gradually be seen to decrease. Also, since a number of elementary regions (225) are being analyzed, it is possible to gain some idea of the average behavior of the correlation fall-off with lack of overlap.

Figures 4-1, 4-2 and 4-3 show the manner in which the correlation between two image segments falls off as the area of overlap is reduced. Figure 4-1 shows the results for the phase correlation algorithm. Figure 4-2 shows the results for the correlation of power-spectrum values, and Figure 4-3 for the moment invariants method.

Results for the moment invariants method are the most uncertain, as indicated by the large error bars (these how the 1σ limits in the correlation estimate in each of the three cases). In conjunction with the results described in the next section, these results can be used to determine the degree of change (size or intensity of targets of interest) which can be detected by these methods.

4.2 TARGET SIZE

As a further experiment, a series of correlation measurements was performed on a pair of identical 32×32 segments taken from one of the SAR image frames, except than an artificial target was embedded in one of the segments. The segment was selected so as to contain no target initially, and for each correlation experiment a target of specified size was embedded. In each case the target intensity was fixed, and was chosen so as to be consistent with the intensity of the stronger targets actually present in the original SAR images. The size of the embedded artificial targets is specified in terms of the area in pixels^2.

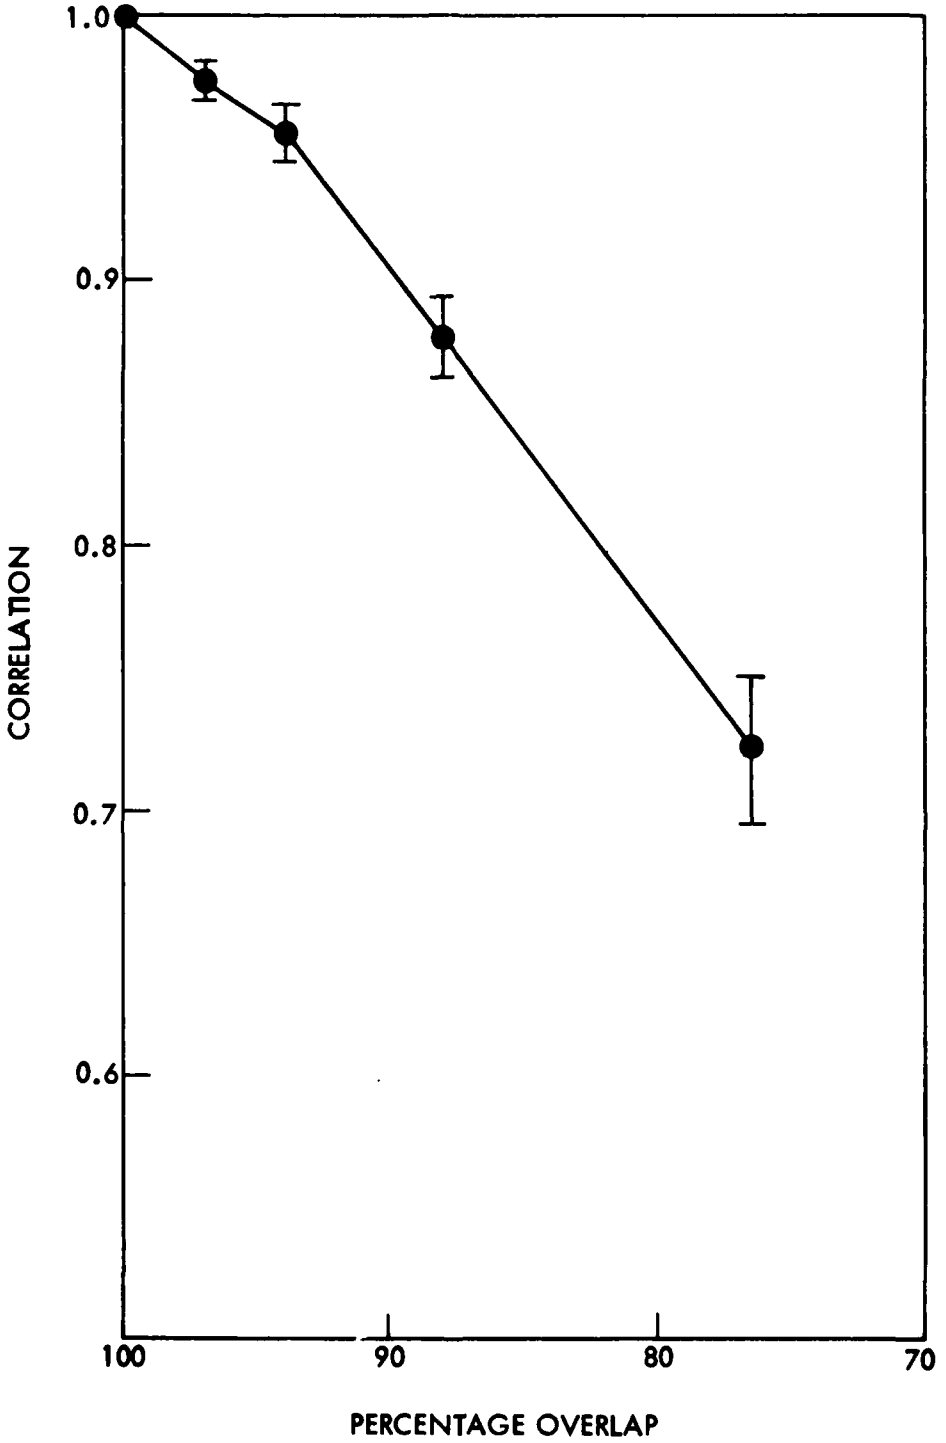


Figure 4-1 Phase-Correlation Method – Cosine Weighting Used; Calculated Correlation Between Two 32 x 32 Image Segments of SAPPHIRE SAR Image, Full-Resolution. 49 Segments Averaged

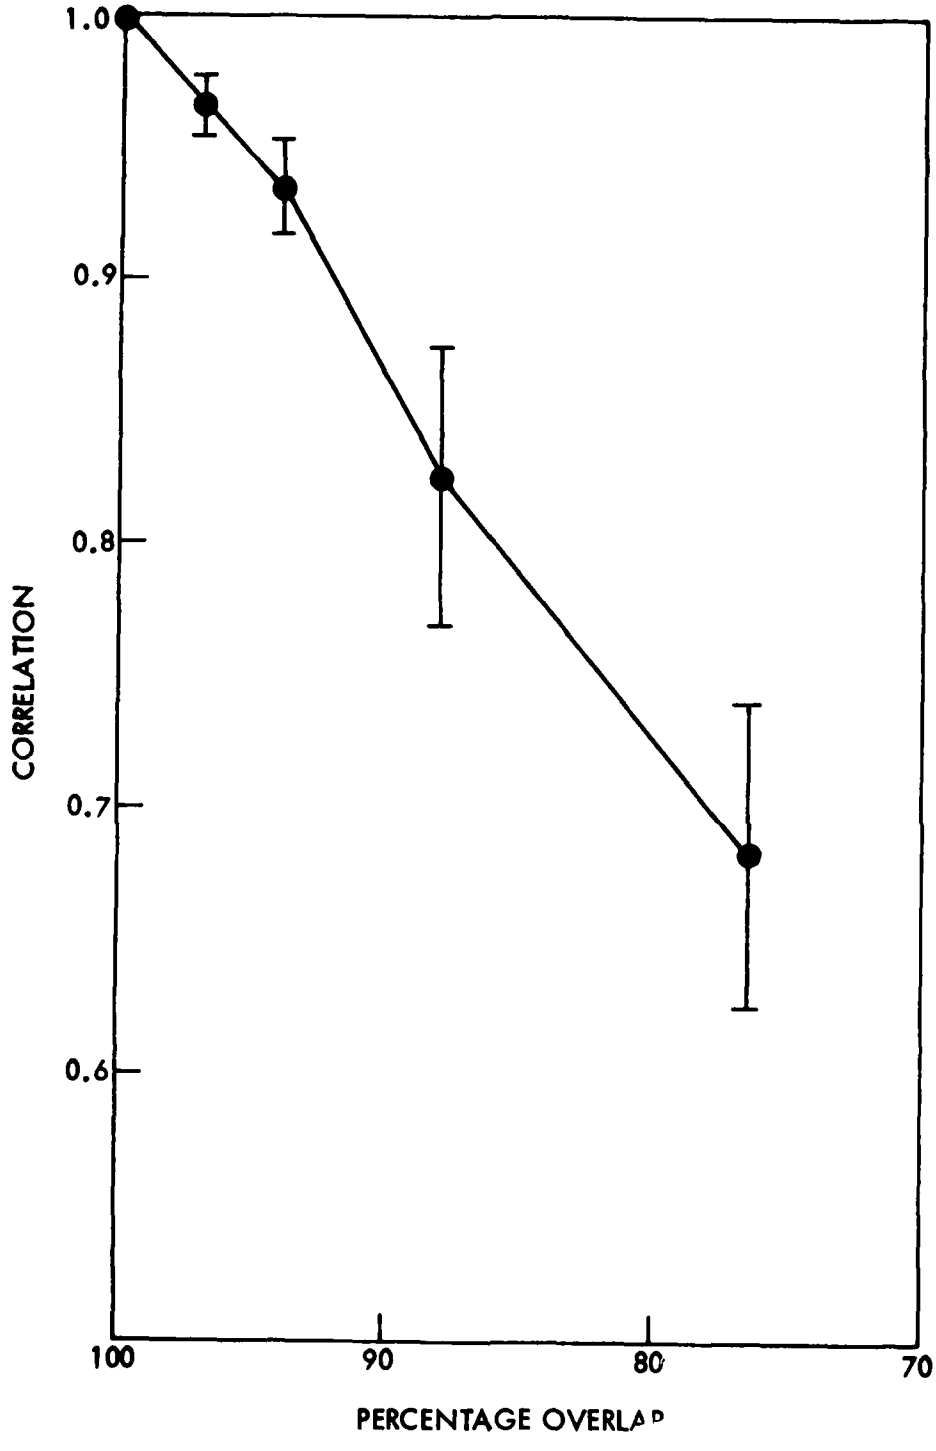


Figure 4-2 Power-Spectrum Correlation Method – Cosine Weighting Used; Calculated Correlation Between Two 32 x 32 Image Segments of SAPPHIRE SAR Image, Full-Resolution. 49 Segments Averaged

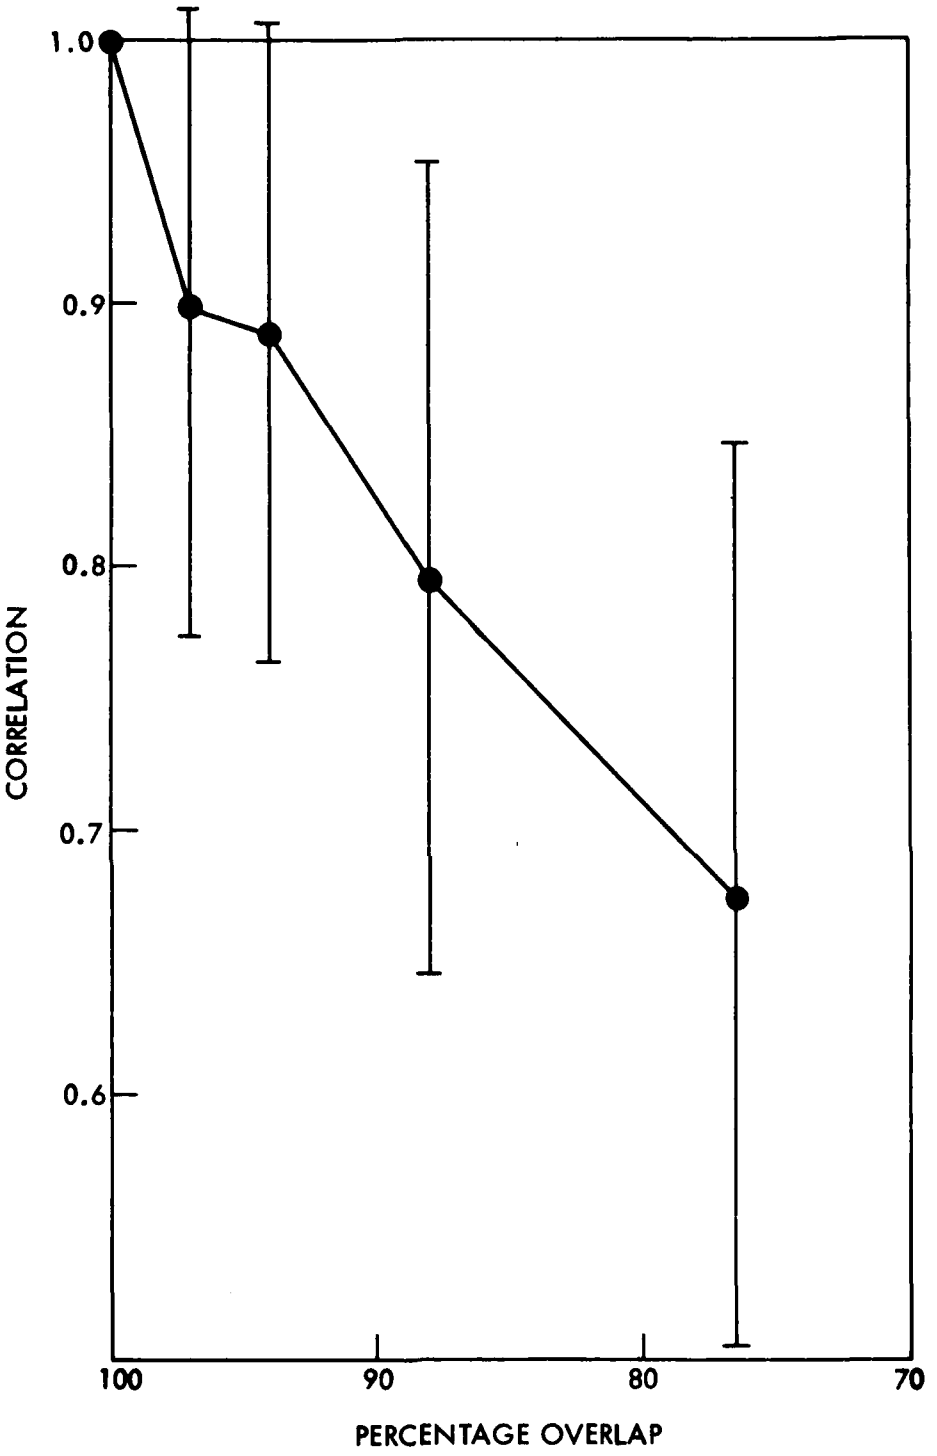


Figure 4-3 Invariants-Correlation Method – Cosine Weighting Used; Calculated Correlation Between Two 32 x 32 Image Segments of SAPPHIRE SAR Image, Full-Resolution. 49 Segments Averaged

Targets of size 1 x 1, 1 x 2, 2 x 2, 4 x 4, 8 x 8 and 16 x 16 were used in the experiments. As the size of the artificial target is increased, the correlation between the two, image segments is reduced in each of the three algorithms. The manner in which this occurs for each is shown by the set of three curves in Figure 4-4, where the correlation coefficient is plotted against the artificial target size (in pixels^2).

It is useful to compare the results of the previous section with the conclusions drawn from the curves of Figure 4-4. A comparison shows that, if significant change is indicated, for example, by the correlation coefficient falling to a value of 0.7 for any of the three algorithms, then a minimum target size for significant change (that is, one that produces the same degree of correlation) is 4 x 4 pixels. This size is rather larger than any of the targets actually present in the available set of SAR images. This consideration is apart from any loss of correlation that two image segments may have because of independent image noise. In the presence of image noise, the target is required to be even larger, in order to indicate significant change.

4.3 IMAGE SIGNAL-TO-NOISE RATIO

The problem of independent noise in each image segment was analyzed in Section 3.4. If significant image noise is present, this becomes the limiting factor in a correlation type of similarity measurement. The minimum target size for any significant change is now increased if both image noise and overlap error are considered.

We can note that, although pixel-averaging was used in the present study as a method of noise reduction, this process may not contribute to greater change detection capability. This may be seen from the fact that the pixel averaging process does not improve the signal/noise ratio by a factor of 2 . If the same size of elementary region is used in both the averaged and the full-resolution cases, targets are now only one-quarter of their forms size and contribute less to change. Use of some fixed-minimum size of elementary region may be required because of the previously discussed overlap problem.

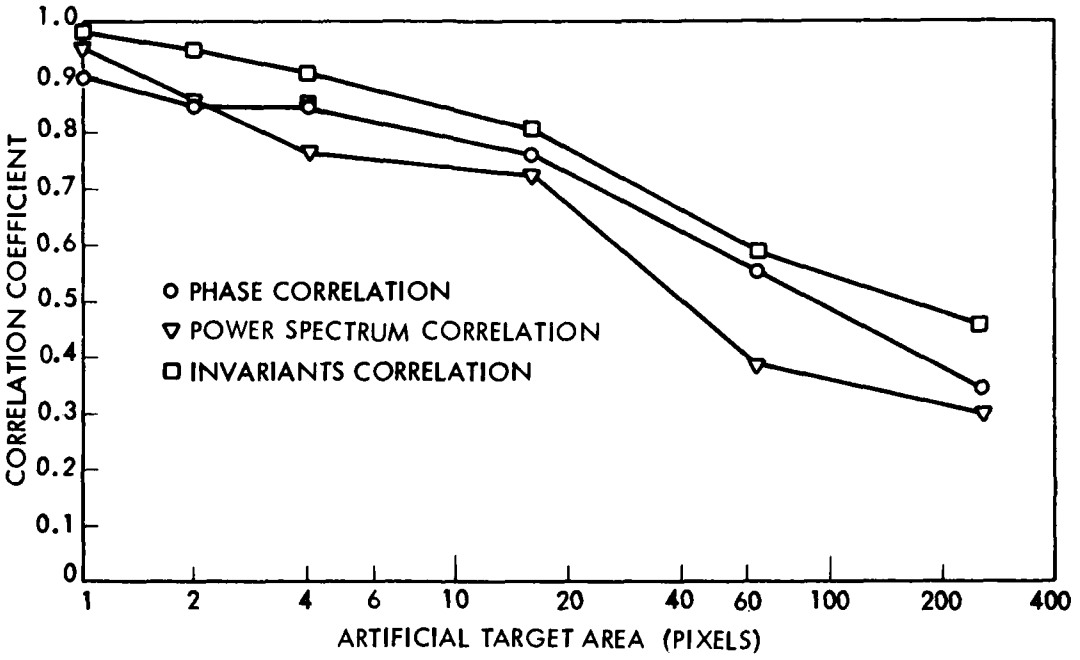


Figure 4-4 Results of Embedding an Artificial Target of Intensity 256. Background Maximum: 230

Section 5

REFERENCES

1. C. D. Kuglin and D. C. Hines, Jr., "The Phase Correlation Image Alignment Method," Proceedings of the IEEE 1975 International Conference on Cybernetics and Society, Sept. 1975, 11. 163-165

2. C. D. Kuglin, Performance of the Phase Correlator in Image Guidance Applications, LMSC-D553148, Lockheed Palo Alto Research Laboratory, Dec. 1976

3. J. J. Pearson, D. C. Hines, Jr., S. Golosman, and C. D. Kuglin, "Video-Rate Image Correlation Processor," Proceedings of the 1977 SPIE Conference, Vol. 119, Dec. 1977

4. M. K. Hu, "Visual Pattern Recognition by Moment Invariants", IRE Transactions on Information Theory, Vol. IT-8, pp. 179-187, February 1962

5. S. Maitra "Moment Invariants", Proceedings of the IEEE, Vol. 67, No. 4 April 1979

MISSION
of
Rome Air Development Center

RADC plans and executes research, development, test and selected acquisition programs in support of Command, Control Communications and Intelligence (C^3I) activities. Technical and engineering support within areas of technical competence is provided to ESD Program Offices (POs) and other ESD elements. The principal technical mission areas are communications, electromagnetic guidance and control, surveillance of ground and aerospace objects, intelligence data collection and handling, information system technology, ionospheric propagation, solid state sciences, microwave physics and electronic reliability, maintainability and compatibility.